arXiv:1404.4560v1 [cs.GT] 17 Apr 2014

A Control Dichotomy for Pure Scoring Rules*

Edith Hemaspaandra
Department of Computer Science
Rochester Institute of Technology
Rochester, NY 14623, USA
www.cs.rit.edu/∼eh

Lane A. Hemaspaandra
Department of Computer Science
University of Rochester
Rochester, NY 14627, USA
www.cs.rochester.edu/u/lane

Henning Schnoor
Institut für Informatik
Christian-Albrechts-Universität zu Kiel
24098 Kiel, Germany
henning.schnoor@email.uni-kiel.de

April 17, 2014

Abstract

Scoring systems are an extremely important class of election systems. A length-m (so-called) scoring vector applies only to m-candidate elections. To handle general elections, one must use a family of vectors, one per length. The most elegant approach to making sure such families are "family-like" is the recently introduced notion of (polynomial-time uniform) pure scoring rules (Betzler and Dorn, 2010), where each scoring vector is obtained from its precursor by adding one new coefficient. We obtain the first dichotomy theorem for pure scoring rules for a control problem. In particular, for constructive control by adding voters (CCAV), we show that CCAV is solvable in polynomial time for k-approval with $k \leq 3$, k-veto with $k \leq 2$, every pure scoring rule in which only the two top-rated candidates gain nonzero scores, and a particular rule that is a "hybrid" of 1-approval and 1-veto. For all other pure scoring rules, CCAV is NP-complete. We also investigate the descriptive richness of different models for defining pure scoring rules, proving how more rule-generation time gives more rules, proving that rationals give more rules than do the natural numbers, and proving that some restrictions previously thought to be "w.l.o.g." in fact do lose generality.

1 Introduction

Elections give rise to a plethora of interesting questions in the social and political sciences, and have been extensively studied from a computer-science point of view in the last two

*Supported in part by NSF grants CCF-{0915792,1101452,1101479} and by COST Action IC1205. Work done in part while H. Schnoor visited RIT supported by an STSM grant of Cost Action IC1205.

decades. We study the *control problem*, in which the chair of an election (ab)uses her power to try to affect the election outcome. In this paper we focus on *constructive control by adding voters* (CCAV), i.e., where the chair tries to make her favorite candidate win by adding voters. Constructive control by adding voters is an extremely important control type, since it occurs in (political) practice very often. A standard example is the "Get-out-the-Vote" efforts of political parties, aimed at (supposed) supporters of those parties. However, we should also mention that, as has been pointed out for example in books by Riker (1986) and Taylor (2005), in modern politics issues of candidate introduction or removal (CCAC/CCDC) have become highly important; the case of Ralph Nader in two recent American elections vividly supports this point.

The computational complexity of CCAV and other forms of control was first studied by Bartholdi, Tovey, and Trick (1992), for plurality and (so-called) Condorcet elections. In this paper, we study the complexity of CCAV for *pure scoring rules*, an attractive class introduced by Betzler and Dorn (2010) that contains many important voting systems. A scoring rule for an election with m candidates is defined by m coefficients $\alpha_1 \geq \alpha_2 \geq \cdots \geq \alpha_m$. Each voter ranks the m candidates from her most favorite to her least favorite; a candidate gains α_i points from being in position i on that voter's ballot.

Well-known examples of families of scoring rules include the following. Borda Count for m candidates uses coefficients $m-1, m-2, \ldots, 1, 0$. k-approval uses coefficients $\underbrace{1, \ldots, 1}_{k}, \underbrace{0, \ldots, 0}_{m-k}$; k-veto uses $\underbrace{1, \ldots, 1}_{m-k}, \underbrace{0, \ldots, 0}_{k}$. Dowdall voting, used for Nauru's parliament, uses $1, \frac{1}{2}, \frac{1}{3}, \ldots, \frac{1}{m}$.

The construction of the scoring vector for a specific number of candidates usually follows a natural pattern, as in the above examples. This leads to the definition of a "pure scoring rule." We discuss the notion of "purity" in detail in this paper. (Basically, it means that at each length we insert one entry into the previous length's vector; all our above examples are pure.)

There is a rich literature on computational aspects of scoring rules, e.g., dichotomy theorems on weighted manipulation (Hemaspaandra and Hemaspaandra, 2007), the possible winner problem (Betzler and Dorn, 2010; Baumeister and Rothe, 2012), and bribery (Faliszewski et al., 2009), as well as results about specific voting systems (Betzler et al., 2011; Davies et al., 2011; Faliszewski et al., 2013).

In this paper, we provide the first complete investigation of the complexity of the unweighted CCAV problem for pure scoring rules. We prove a dichotomy theorem that gives a complete complexity-theoretic classification of that control problem for pure scoring rules. Our result is as follows.

It turns out that there are only 4 types of pure scoring rules for which CCAV is solvable in polynomial time:

1. k-approval for $k \leq 3$,

2. k-veto for $k \leq 2$,

3. every pure scoring rule in which only the two top-rated candidates receive a nonzero score,

4. a particular rule which is a "hybrid" of 1-approval and 1-veto: each voter awards her favorite candidate 1 point, and her least favorite candidate -1 point.

For every pure scoring rule that is not one of the above election systems, CCAV is NP-complete. The last rule mentioned above is particularly interesting for two reasons: First, it was the only one for which the complexity of the possible winner problem was left open in Betzler and Dorn (2010). Second, it is the only one for which polynomial-time solvability depends on the actual coefficients, and not only on the $<$-order of the values: While this election system is equal to the rule generated by coefficients $2, 1, \ldots, 1, 0$, the rule generated by using coefficients $3, 1, \ldots, 1, 0$ leads to an NP-complete CCAV-problem.

A key point in the proof of any dichotomy theorem is to cover all relevant cases. We thus base our dichotomy result on a study of the descriptional richness of definitions of scoring protocols. We examine how variations of key parameters as the abovementioned purity requirement, the complexity allowed to compute the coefficients, and the universe from which the coefficients may be chosen affect the set of election systems that can be represented. Interestingly, we discover that assumptions made previously in the literature, which were believed to be without loss of generality, in fact restrict the rules that can be expressed. Taking these results into account, we introduce a flexible purity condition that is more robust with respect to the abovementioned variations. We show that the new notion strictly generalizes pure scoring rules with integer scores and coincides with pure scoring rules with rational coefficients.

The paper is structured as follows. In Section 2, we give necessary preliminaries about the class of election systems that we study. In Section 3, we study the descriptive richness of definitions of positional scoring rules. Section 4 contains our main complexity result. All proofs from Section 3 and some technical proofs from Section 4 have been deferred to the appendix.

2 Preliminaries

As is standard, given an m-component (so-called "scoring") vector $\alpha = (\alpha_1, \ldots, \alpha_m)$ such that $\alpha_1 \geq \ldots \geq \alpha_m$, we define a so-called (m-candidate) "scoring system" election based on this by, for each voter (the voters vote by strict, linear orders over the candidates), giving α_i points to the voter's ith-most-favorite candidate. Whichever candidate(s) get the highest number of points, summed over all voters, are the winner(s). We refer to the α_is as "coefficients."

We will use the following very pure definition of pure scoring rules, which removes some of the additional assumptions that have been used in earlier work. Indeed, earlier work asserted that (at least some of) these assumptions were not restrictions; we will look at that issue anew below.

An election system $\mathcal{E}$ is a $\mathcal{T}$-GSR (generalized scoring rule) if there is a function f that on each input 0^m, $m \geq 1$, outputs an m-component scoring vector $\alpha^m = (\alpha_1^m, \ldots, \alpha_m^m)$ such that $\alpha_1^m \geq \ldots \geq \alpha_m^m$ and each α_j^i belongs to $\mathcal{T}$, and for each m, the winner set under $\mathcal{E}$ is exactly the winner set given by the scoring system using the scoring vector α^m. The notation 0^m denotes a string of m "0"s. Using this as the argument ensures that the generator's computation time is measured as a function of m, since its input length is exactly m. Since the votes are of size at least m each, this provides the natural, fair approach to framing FP-uniformity (as we will do below). We call f a generator for $\mathcal{E}$. While one can consider election systems based on an f that is not computable, in practice we want there to be an efficient algorithm computing f. This is expressed with different *uniformity* conditions. If $\mathcal{E}$ is a $\mathcal{T}$-GSR via some generator f that can be computed in a complexity class $\mathcal{F}$, then $\mathcal{E}$ is an $\mathcal{F}$-uniform-$\mathcal{T}$-GSR and f is an $\mathcal{F}$-uniform $\mathcal{T}$-generator. The values of $\mathcal{T}$ we will be interested in are the naturals $\mathbb{N}$, the nonnegative rationals $\mathbb{Q}_{\geq 0}$, the rationals $\mathbb{Q}$, and the integers $\mathbb{Z}$. Our most important value of $\mathcal{F}$ will be the polynomial-time functions, FP. When we do not state "nonuniform" or some specific uniformity, we always mean FP-uniform. When we put a name in boldface, it indicates all the elections that can be generated by a generator of the named sort, e.g., **FP-uniform-$\mathbb{Z}$-GSR** is the class of all polynomial-time-uniform generalized scoring rules with integer coefficients, a class first defined and discussed by Hemaspaandra and Hemaspaandra (2007).

A $\mathcal{T}$-GSR (of whatever uniformity) is a $\mathcal{T}$-PSR (pure scoring rule) if it has a generator (of the same uniformity) f satisfying the following "purity" constraint: For each $m \geq 2$, there is a component of α^m that when deleted leaves exactly the vector α^{m-1}.

Throughout the paper we use the following observation, which notes that two scoring vectors, after being "normalized," differ if and only if they are capturing distinct election systems: An m-position scoring vector $\alpha = (\alpha_1, \ldots, \alpha_m)$ (over $\mathbb{N}$) is normalized if $\alpha_m = 0$ and the greatest common divisor of its nonzero α_i's is 1. Normalizing a given scoring vector (over $\mathbb{N}$ or $\mathbb{Z}$) is easily achievable in polynomial time: subtract α_m from each coefficient and then divide each coefficient by the gcd (computed using Euclid's gcd algorithm) of the nonzero thus-altered coefficients. The normalization of a scoring vector over $\mathbb{Q}$ (resp., $\mathbb{Q}_{\geq 0}$) is done by multiplying through by the lcm of the denominators of the nonzero coefficients, and then viewing that as a vector over $\mathbb{Z}$ (resp., $\mathbb{N}$) and normalizing it as above.

Proposition 2.1. *Let $m \geq 1$ and let α and α' be m-position scoring vectors over $\mathcal{T} \subseteq \mathbb{Q}$. Then α and α' have the same winner sets on each m-candidate election if and only if α and α' both have the same normalized version.*

The if direction basically follows from Observation 2.2 of Hemaspaandra and Hemaspaandra (2007), as noted by Betzler and Dorn (2010). The only if follows by giving a construction that for any two unequal normalized scoring vectors constructs a vote set on which their winner sets differ. The construction works by "aligning" the vectors by multiplying each so that their first coefficients are equal, and then using a padding construction to ensure that only two candidates are crucial and that the winner sets can be distinguished by appropriately exploiting the first position at which the aligned vectors differ.

Generators f_1 and f_2 are *equivalent* if they generate the same election system. Due to Proposition 2.1, this is the case if and only if, for each length m, the *normalized* scoring vectors generated by f_1 and f_2 for m candidates are identical.

3 Descriptive Richness and PSRs

We now examine how amount of time used to generate PSRs, and the universe the PSR's coefficients are drawn from, affect the family of election rules that can be obtained. We also look at whether such a seemingly innocuous and standard assumption as having the last coefficient always being zero in fact loses generality; we'll see that it does lose generality, but in a way that can in part be papered over.

3.1 Generation Time Gives Descriptional Richness

Does more generation time give a richer class of pure scoring rules? A very tight time hierarchy can be achieved by a legal form of cheating. In particular, consider any (nice) time class that can for some $m \geq 3$ generate a scoring vector over $\mathbb{N}$ of the form $(\alpha_1, \underbrace{1, \ldots, 1}_{m-2}, 0)$ such that α_1 is so big that some other time class cannot generate this scoring vector (for example, because it simply doesn't have enough time to write down enough bits to get a number as large as α_1). Our vector is normalized, and by Proposition 2.1 this already is enough to allow us to argue that the two time classes differ in their winner sets. (This proof works for GSRs. And it works for PSRs if the "nice"ness of the class allows it to obey purity while employing the above approach—hardly an onerous requirement.) However, that claim simply uses the fact that more time can write more bits. A truly fair and far more interesting separation would show that one can with more time obtain more pure scoring rules in a way that does *not* depend on using coefficient lengths that simply cannot be produced by the weaker time class. In the dream case, all coefficients in fact would simply be 0 or 1, so there are no long coefficients in play at all. We in fact have achieved such a hierarchy theorem. Full time constructibility is a standard notion that most natural time functions satisfy and, as is standard, when we speak of a time function $T(m)$ by convention that is shorthand for $\max(m+1, \lceil T(m) \rceil)$ (Hopcroft and Ullman, 1979). FDTIME$[g(\cdot)]$ denotes the functions computable in the given amount of deterministic time.

Theorem 3.1. *If $T_2(m)$ is a fully time-constructible function and $\limsup_{n\to\infty} \frac{T_1(m)\log T_1(m)}{T_2(m)} = 0$, then there is an election rule in* **FDTIME$[T_2(m)]$-*uniform*-$\{0,1\}$-PSR** *that is not in* **FDTIME$[T_1(m)]$-*uniform*-$\mathbb{Q}$-GSR**.

The log factor here is not surprising; this is the standard overhead it takes for a 2-tape Turing machine to simulate a multitape TM. What is surprising is that no additional factor of m is needed. Why might we expect such a factor? (We caution that the rest of this paragraph is intended mostly for those having familiarity with the diagonalization techniques used to prove hierarchy theorems.) In the context of a diagonalization construction (which

is the basic technique used in the proof of Theorem 3.1), one might expect to (all counting against the overall time T_2 limit) at vector-length m have to "recreate" all the shorter vectors used in earlier diagonalizations to ensure that the length-m and length-$(m-1)$ vectors are related in a "pure" way. We however sidestep the need for that overhead by a "purity"-inducing trick: For each odd m our scoring vector will be of the form $1^{\lfloor m/2 \rfloor}0^{\lfloor m/2 \rfloor+1}$, and at each even length, we purely extend that to whichever of $1^{\lfloor m/2 \rfloor+1}0^{\lfloor m/2 \rfloor+1}$ or $1^{\lfloor m/2 \rfloor}0^{\lfloor m/2 \rfloor+2}$ diagonalizes against the T_1-time machine that is being diagonalized against (if this is an m when we have time to so diagonalize). Briefly, we lurch back to fixed, safe way-stations at every second length, and this removes the need to recompute our own history.

3.2 Coefficient Richness Gives Descriptional Richness

The richness and structure of the coefficient set for PSRs affects how broad a class of election rules can be captured, as shown by the following result. (Of course, trivially $\mathbb{N}$-**PSR** $\subseteq \mathbb{Q}_{\geq \mathbf{0}}$-**PSR** $\cap$ $\mathbb{Z}$-**PSR** $\subseteq \mathbb{Q}_{\geq \mathbf{0}}$-**PSR** $\cup$ $\mathbb{Z}$-**PSR** $\subseteq \mathbb{Q}$-**PSR.**)

Theorem 3.2. $\mathbb{Q}_{\geq \mathbf{0}}$-**PSR** $\not\subseteq$ ***nonuniform***-$\mathbb{Z}$-**PSR**, $\mathbb{Z}$-**PSR** $\not\subseteq$ ***nonuniform***-$\mathbb{Q}_{\geq \mathbf{0}}$-**PSR**, *and* $\mathbb{Q}$-**PSR** $\not\subseteq$ ***nonuniform***-$\mathbb{Q}_{\geq \mathbf{0}}$-**PSR** $\cup$ ***nonuniform***-$\mathbb{Z}$-**PSR**.

So for example, in pure scoring rules, (polynomial-time uniformly) using positive rationals cannot be simulated by naturals or integers, even nonuniformly. And in pure scoring rules, (polynomial-time uniformly) using integers cannot be simulated by naturals or positive rationals, even nonuniformly. One might think these claims are impossible, and that by normalizing one can go back and forth, but it is precisely the purity requirement that is making that sort of manipulation impossible—there is a price to purity, and it is showing itself here. (The final part of the theorem does not follow automatically from the first two parts plus the trivial observation before the theorem; just the weaker variant of that part in which $\cup$ is replaced with $\cap$ follows from those.) Note that enlarging the universe does not necessarily lead to a larger class of election systems: For example, requiring that coefficients are odd natural numbers gives the same set of election systems as merely requiring them to be natural numbers.

We mention a more flexible and highly attractive notion of purity that erases the differences just discussed. $\mathcal{T}$-FPSRs (flexible pure scoring rules), of whatever uniformity or nonuniformity, will be defined exactly the same way as $\mathcal{T}$-PSRs were defined, except the purity condition is changed to: For each $m \geq 2$, there is a component of α^m whose removal gives a scoring vector equivalent to α^{m-1}. Due to Proposition 2.1, this means that the relevant scoring vectors have the same normalization. We call such generators *flexible*. For *this* notion we have for the nonuniform case and the FP-uniform case (and most other nice cases), the following equality.

Theorem 3.3. $\mathbb{N}$-**FPSR** $=$ $\mathbb{Q}$-**FPSR** $=$ $\mathbb{Q}_{\geq \mathbf{0}}$-**FPSR** $=$ $\mathbb{Z}$-**FPSR** $=$ $\mathbb{Q}$-**PSR**.

3.3 Having Smallest Coefficient of Zero Loses Generality—Slightly

Betzler and Dorn (2010) in their definition of scoring rules require that at each m, we have $\alpha_m^m = 0$ (let us call this condition norm-0), and comment that this is not a restriction. We note that there are PSRs that cannot be generated by any pure scoring rule meeting that constraint. What is at issue here is a bit subtle: At each fixed length, the restriction is innocuous. But in the context of families that are bound by the purity constraint, the restriction loses generality. On the other hand, we will also note that each pure scoring rule has a generator that is "close" to meeting that constraint—it meets it at all but finitely many lengths.

Betzler and Dorn (2010) also have a "gcd is 1" condition (although the phrasing is not crystal clear as to whether the gcd constraint applies to the union of all nonzero coefficients that occur over all lengths, or whether it must in fact apply separately at each length; the latter (which we call norm-gcd) would block the vector $(2, 0)$ but the former would not if the next step were, for example, $(3, 2, 0)$). However, the gcd issues also can be made to go away "almost everywhere." In particular, we have established the following.

Theorem 3.4. *There is an* FP*-uniform* PSR *that is not generated even by any nonuniform* $\text{PSR}_{\textit{norm-0}}$ *generator. On the other hand, every* FP*-uniform (respectively, nonuniform)* PSR *is generated by a* FP*-uniform (respectively, nonuniform)* PSR *generator that for all but a finite number of* m *has the property that the last coefficient,* $\alpha_{m,m}$*, is zero and the gcd of the nonzero coefficients in the length-*m *vector is one.*

Does the second sentence of the theorem imply that each PSR has all its vectors the same at each length (except for a finite number of exceptional lengths) as the vectors of some PSR that satisfies norm-0 and norm-gcd? The answer is actually "no." The somewhat subtle issue at play is that PSRs can generate vectors that no generator satisfying norm-0 and norm-gcd can ever generate, such as the family $(3, 2, \ldots, 2, 0)$. So one should not from our above theorem claim that it follows from the main dichotomy *theorem* of Betzler and Dorn (2010), as completed by Baumeister and Rothe (2012), that we can read off the complexity (of the possible winner problem) even in our slightly more flexible case. However, our above theorem does—in light of the actual *proof case decomposition* used in those papers (which is based on issues such as whether one has an unbounded number of positions that differ and so on) and some additional argumentation to connect to that and in particular to note that Betzler and Dorn (and Baumeister and Rothe) are in effect quietly covering well even those cases that do not satisfy gcd constraints—connect so well to their work that each of our cases is settled by their proofs.

4 A Control Dichotomy for PSRs

We study the following problem for an election system $\mathcal{E}$: When R is a set of registered voters, is there some subset of the unregistered voters U of size at most k that we can add to the election to ensure that p is the winner?

Definition 4.1. *Let $\mathcal{E}$ be an election system. The* constructive control problem for $\mathcal{E}$ by adding voters*, $\mathcal{E}$-CCAV, is the following problem: Given two multisets sets of votes R and U, a candidate p and a number k, is there a set $A \subseteq U$ with $\|A\| \leq k$ such that p is a winner of the election if the votes in the multiset $R \cup A$ are evaluated using the system $\mathcal{E}$?*

We often use a generator, f, as a shorthand for the election system (scoring rule family) it generates, e.g., we write f-CCAV. For a generator f, we use $\alpha^{f,m} = (\alpha_1^{f,m}, \ldots, \alpha_m^{f,m})$ to denote the scoring vector generated by f for m candidates and its individual coefficients. To simplify presentation, we only consider FP-uniform generators. However our results continue to hold as long as we can solve the following question in polynomial time: Given m, i, and j in *unary*, does $\alpha_i^m > \alpha_j^m$ hold, where $f(m) = (\alpha_1^m, \ldots, \alpha_m^m)$?

Our main result is a complexity dichotomy for f-CCAV when f is an FP-uniform pure $\mathbb{Q}$-generator (or, equivalently due to Theorem 3.3, a FP-uniform flexible $\mathbb{N}$-generator). Recall that equivalent generators result in the same election system, hence, due to Proposition 2.1, Theorem 4.2 implies polynomial-time results for all generators with the same normalization as one below. We state our main result.

Theorem 4.2. *Let f be a* FP*-uniform pure $\mathbb{Q}$-generator. Then f-CCAV is solvable in polynomial time if f is equivalent to one of the following generators:*

- $f_1 = (1, 1, 1, 0, \ldots, 0)$ *(this generates* 3*-approval),*
- $f_2 = (1, \ldots, 1, 0)$ *or* $f_3 = (1, \ldots, 1, 0, 0)$ *(1/2-veto),*
- *for some* $\alpha \geq \beta$*,* $f_4 = (\alpha, \beta, 0, \ldots, 0)$*,*
- $f_5 = (2, 1, \ldots, 1, 0)$*.*

In all other cases, f-CCAV is NP-complete.

Note that 1- and 2-approval are covered by the generator f_4. The remainder of the paper contains the proof of Theorem 4.2: Section 4.1 contains the algorithms for all polynomial-time solvable cases, Section 4.2 contains our hardness results (with some proofs deferred to the appendix), and the proof of Theorem 4.2 follows in Section 4.3.

4.1 Polynomial-Time Results

The following result is proven by Lin (2012).

Theorem 4.3. *$\mathcal{E}$-CCAV is solvable in polynomial time if $\mathcal{E}$ is k-approval with $k \leq 3$ or k-veto with $k \leq 2$.*

Due to Proposition 2.1, Theorem 4.3 implies that CCAV remains polynomial-time solvable for "scaled" versions of k-approval with $k \leq 3$ or k-veto with $k \leq 2$, i.e., generators of the form $(\alpha, \beta, \gamma, \delta, \ldots, \delta)$ with $\beta, \gamma \in \{\alpha, \delta\}$ or $(\alpha, \ldots, \alpha, \beta, \gamma)$ with $\beta \in \{\alpha, \gamma\}$. We now look at a generalization of 2-approval: Voters approve of 2 candidates, and can distinguish

between their first and second choice. CCAV for this generalization remains efficiently solvable. In contrast, Theorem 4.11 shows that our control problem for the corresponding generalization of 3-approval is NP-hard.

Theorem 4.4. *Let $\alpha \geq \beta$ be fixed. Then f-CCAV is polynomial-time solvable for $f = (\alpha, \beta, 0, \ldots, 0)$.*

Proof. Due to Theorem 4.3, assume $\alpha > \beta$. Let ℓ be such that $\ell(\alpha - \beta) \geq \beta$. Let an instance with candidates C, favorite candidate p, registered voters R, and potential voters U be given; let k be the number of voters that can be added. Assume $\ell \leq k$, otherwise brute-force. Let V_1 (V_2) be the set of voters in U that put p in the first (second) spot. Clearly, we add voters only from $V_1 \cup V_2$. Assume w.l.o.g. that two voters who vote identically in the first two positions also rank the remaining candidates identically. In particular, two voters in V_1 (V_2) are different if and only if they vote different candidates in the second (first) place. We use the following facts.

Fact 1. *For $i \in \{1, 2\}$, given a set $S \subseteq V_1 \cup V_2$, it can be checked in polynomial time whether S can be extended, by adding at most $k - \|S\|$ voters from V_i to make p win.*

Fact 2. *To make p win, it is never better to add ℓ voters from V_2 than adding ℓ pairwise different voters from V_1.*

Due to Fact 2, we do not have to consider solutions that use ℓ or more voters from V_2 and leave ℓ or more distinct voters from V_1 unused. So if a solution exists, we can find one using fewer than ℓ voters from V_2, or leaving fewer than ℓ pairwise different voters from V_1 unused. For both cases we will test whether there is a corresponding solution.

We start with the first case. Since ℓ is constant, we can test every subset $S \subseteq V_2$ with $\|S\| < \ell$. For each of these S, we use Fact 1 to check in polynomial time whether S can be extended to a solution by adding voters only from V_1.

For the second case, we determine in polynomial time whether there is a solution that does not leave ℓ pairwise different voters from V_1 unused as follows: We encode the choice of the unused voters from V_1 as a function $u\colon C \to \{0, \ldots, \|V_1\|\}$ that states, for each candidate c, the number of unused V_1-voters placing c second. Since we look for solutions satisfying the second case, we only consider functions u for which $u(c) > 0$ for at most $\ell - 1$ candidates.

Thus we can brute-force over all of these possibilities as follows:

1. In the outer loop, we test every subset $S \subseteq C$ with $\|S\| \leq \ell - 1$. Since ℓ is a constant, there are only polynomially many of these sets.

2. For each such S, we test all functions u as above for which $u(c) > 0$ holds iff $c \in S$. Such a function can be regarded as a function $u\colon S \to \{1, \ldots, \|V_1\|\}$. There are $\|V_1\|^{\|S\|} \leq \|V_1\|^{\ell-1}$ many of these functions. Since $\|V_1\|$ is bound by the input size and ℓ is a constant, this number is polynomial in the input size.

So we can polynomially go through all possibilities of potentially unused V_1-voters, which is the same as going through all possible sets S' of *used* V_1-voters. For each of these sets S', we again use Fact 1 to check in polynomial time whether S' can be extended to a solution. This completes the proof. □

Our final polynomial-time case is the generator $(2, 1, \ldots, 1, 0)$. Here every voter "approves" one candidate and "vetoes" another. This case is interesting for two reasons. First, it is the only case where the algorithm depends on the coefficients itself, as opposed to their $>$-order. Namely, for all $\alpha > \beta > 0$ with $\alpha \neq 2\beta$, f-CCAV with $f = (\alpha, \beta, \ldots, \beta, 0)$ is NP-complete (Theorem 4.12.2). Second, this case was the only one left open in Betzler and Dorn's (2010) possible winner dichotomy; the question was eventually settled by Baumeister and Rothe (2012), who proved NP-completeness.

Theorem 4.5. *f-CCAV is solvable in polynomial time for the generator $f = (2, 1, \ldots, 1, 0)$.*

Proof. Let C, R, and U be the set of candidates, registered voters, and unregistered voters, p the preferred candidate, and k the number of voters we can add. We add no voter voting p last, and it is never better to add a voter voting p second than to add one voting p first. So we first add all voters from U that place p in the first position. If there are more than k of these voters, we choose the ones to add with the obvious greedy strategy that always picks, among all available votes of the form $p > \cdots > c$, the one where c currently has the highest score. After this preprocessing, all relevant voters in U vote $c_1 > \cdots > c_2$ with $p \notin \{c_1, c_2\}$. To simplify presentation, we use Proposition 2.1 and consider f as the generator $(1, 0, \ldots, 0, -1)$. Then the score of p is determined by the votes in R.

We reduce the problem to min-cost (network) flow, which can be solved in polynomial time. Let $S = \sum_{c \in C - \{p\}} \mathit{score}(c)$. We use the following nodes and edges:

- For each $c \in C - \{p\}$, there is a node c, additionally, there are source and target nodes s and t.
- There is an edge from candidate c_1 to candidate c_2 with cost 1 and with capacity equal to the number of voters in U voting $c_2 > \cdots > c_1$.
- For each candidate-node c, there is an edge from s to c with cost 0 and capacity $\mathit{score}(c)$ and an edge from c to t with cost 0 and capacity $\mathit{score}(p)$.

Now p can be made winner with at most k additional voters if and only if there is a flow from s to t with value S and cost at most k: Clearly, network flows with cost at most k correspond to subsets of U with size at most k, and using an edge (c_1, c_2) r times corresponds to adding r voters voting $c_2 > \cdots > c_1$, since this vote transfers one point from c_1 to c_2. The capacity of the outgoing edges of s ensure that each candidate initially gets the correct number of points (since S points must be distributed), the edges to t ensure that in the end, no candidate may have more points than p. □

The above results cover all polynomial-time cases of Theorem 4.2. We now turn to the NP-complete cases.

4.2 Hardness Results

We use the standard NP-complete problem 3DM (3-dimensional matching).

Definition 4.6. *3DM is defined as follows:*

Input	*Pairwise disjoint sets X, Y, and Z with $\|X\| = \|Y\| = \|Z\|$, and a set $M \subseteq X \times Y \times Z$.*
Question	*Is there a set $C \subseteq M$ with $\|C\| = \|X\|$ that covers X, Y, and Z?*

We say that C covers X (resp., Y, Z) if every element from X (resp., Y, Z) appears in a tuple of C. Since X, Y, and Z are pairwise disjoint, in this case every element from X (Y, Z) appears in the first (second, third) component of a tuple from C. Since $\|X\| = \|Y\| = \|Z\|$, a set $C \subseteq M$ with $\|C\| = \|X\|$ covers X (Y, Z) if and only if no two tuples from C agree in the first (second, third) component. A set C covering X, Y, and Z is called a *cover*.

4.2.1 Constructing Elections

In our hardness proofs, we often need to set up the registered voters to ensure specific scores for the candidates. The following lemma shows that, if there is a "dummy" candidate to whom any surplus points can be "shifted," we can obtain every set of relative scores that can be expressed as a polynomial-size linear combination of the coefficients in the scoring vector.

Lemma 4.7. *Given a scoring vector $(\alpha_1, \ldots, \alpha_m)$, and for each $c \in \{1, \ldots, m-1\}$, numbers $a_1^c, \ldots, a_m^c$ in signed unary, and a number k in unary, we can compute, in polynomial time, votes such that the scores of the candidates when evaluating these votes according to the scoring vector $(\alpha_1, \ldots, \alpha_m)$ are as follows: There is some o such that for each $c \in \{1, \ldots, m-1\}$, $score(c) = o + \sum_{i=1}^{m} a_i^c \alpha_i$, and $score(c) > score(m) + k\alpha_1$.*

The value o in Lemma 4.7 is the common offset for all relevant scores. The actual value of o is irrelevant, since the winner of the election is determined by the relative scores. The value k is given so that the computed votes ensure that the dummy candidate m cannot win the election with the addition of at most k voters.

4.2.2 "Many" Different Coefficients

We now show that the CCAV-problem is NP-complete for generators using "many" different coefficients. Consider any generator f using (at least) 7 different coefficients for some length m. Then with $\alpha^{f,m} = (\alpha_1^{f,m}, \alpha_2^{f,m}, \alpha_3^{f,m}, \alpha_4^{f,m}, \ldots, \alpha_{m-2}^{f,m}, \alpha_{m-1}^{f,m}, \alpha_m^{f,m})$ we know that $\alpha_4^{f,m} > \alpha_{m-2}^{f,m}$. This condition in fact suffices for the CCAV problem to be NP-hard; the result applies to, e.g., Borda, 3-veto, and 4-approval (the latter two use just two different coefficients, but satisfy $\alpha_4^{f,m} > \alpha_{m-2}^{f,m}$ for $m \geq 7$).

For 4-approval or 3-veto, NP-hardness can be proven by positioning the elements of M from a 3DM-instance, along with p, in the 4 top positions of an unregistered 4-approval vote or (without p) in the last 3 positions of an unregistered 3-veto vote. In our cases, we can

always "simulate" one of these systems: If $\alpha_4^{f,m} > \alpha_{m-2}^{f,m}$, then being ranked in one of the first 4 positions is strictly better than being ranked in one of the last 3 positions. Roughly speaking, if "many" intermediate coefficients are larger than the last 3, then the last 3 are the "exception," and we can use them to "simulate" 3-veto. On the other hand, if "many" intermediate coefficients are smaller than the first 4, then the first 4 are the "exception" and we "simulate" 4-approval.[1] NP-hardness for both 3-veto and 4-approval is proved by Lin (2012); however we use a direct reduction from 3DM in our generalization.

We start with the "simulation" of 3-veto. The statement of the following result is a bit unusual. It indeed gives a reduction for generators meeting the condition $\alpha_{3k+1}^{f,m} > \alpha_{m-2}^{f,m}$ for all m. But beyond that the function g gives what we call a "partial" reduction from 3DM to f-CCAV for f's that meet the condition for some values of m. In the proof, the size of the 3DM instance is artificially enlarged to ensure that this "partial reduction" meets an analogue counterpart in such a way that for every generator f that satisfies $\alpha_4^{f,m} > \alpha_{m-2}^{f,m}$ for *some* m, we know that for *each* large enough m, one of the two reductions can be applied. (Appendix Section B.3 has more on this.)

Theorem 4.8. *Let f be an* FP*-uniform $\mathbb{Q}$-generator. Then there exists an* FP*-computable function g such that*

- *g takes as input an instance I_{3DM} of 3DM and produces an instance I_{CCAV} of f-CCAV with $m = 6k$ candidates, where $k = \|X\| = \|Y\| = \|Z\|$.*
- *If $\alpha_{3k+1}^{f,m} > \alpha_{m-2}^{f,m}$, then: I_{3DM} is a positive instance of 3DM iff I_{CCAV} is a positive instance of f-CCAV.*

Proof. We write α_i for $\alpha_i^{f,m}$. W.l.o.g., let $X = \{s_1, \ldots, s_k\}$, $Y = \{s_{k+1}, \ldots, s_{2k}\}$, and $Z = \{s_{2k+1}, \ldots, s_{3k}\}$. We use the following candidates:

- Each $s_i \in \{s_1, \ldots, s_{3k}\}$ is a candidate.
- p is the preferred candidate.
- There are dummy candidates $d_1, \ldots, d_{m-3k-1}$. We assume there are at least 3 dummy candidates, i.e., $k \geq 2$.

We now use Lemma 4.7 to construct the set R of registered voters such that the scores of the candidates are as follows. (In the following, we "normalize" the scores of all candidates using the score of p as a base. So we pretend that the number o from the application of Lemma 4.7 is zero in order to simplify the presentation, clearly the absolute points of all candidates must be positive and are shifted by the actual number o from the lemma.)

- $score(p) = 0$.

[1] For generators where both cases apply such as $f = (2, 2, 2, 2, 1, 1, 1, \ldots, 1, 0, 0, 0)$, either reduction works.

- $score(s_i) = k\alpha_1 - (k-1)\alpha_{1+i} - \alpha_{m-r(i)}$, where $r(i) = 2, 1, 0$ depending on whether $s_i \in X, Y, Z$, respectively. (So $\alpha_{m-r(i)}$ is exactly the amount of points that s_i gains from a vote that "vetoes" s_i, see below)

- $score(d_i) < -k\alpha_1$ for $i \in \{1, \ldots, d_{m-3k-1}\}$.

Let $M \subseteq X \times Y \times Z$ be the set from $I_{3\text{DM}}$. For each $(x, y, z) = (s_h, s_i, s_j) \in M$ (so clearly $h < i < j$), we add an available voter to U voting as follows:

$p > s_1 > \cdots > s_{h-1} > d_1 > s_{h+1} > \cdots > s_{i-1} >$
$d_2 > s_{i+1} > \cdots > s_{j-1} > d_3 > s_{j+1} > \cdots > s_{3k} >$
$d_4 > \cdots > d_{m-3k-1} > x > y > z.$

We say that such a vote *vetoes* the candidates x, y, and z, and identify elements of M and the corresponding votes.

We show that the reduction is correct. First assume that the instance of 3DM is positive, and let $C \subseteq M$ be the cover with $\|C\| = k$. We add the voters corresponding to the elements of C in the obvious way and show that p indeed wins the resulting election.

To see this, it suffices to show that p has at least as many points as each candidate s_i, since by construction, the dummy candidates cannot win the election with adding at most k votes. So let $i \in \{1, \ldots, 3k\}$. The final score for p and s_i are as follows:

- p gains α_1 points in each of the k additional votes, so p ends up with exactly $k\alpha_1$ points.

- s_i gains $(k-1)\alpha_{1+i}$ points from the $(k-1)$ votes corresponding to elements $(x, y, z) \in C$ with $s_i \notin \{x, y, z\}$, and $\alpha_{m-r(i)}$ points from the single vote vetoing s_i. So s_i ends up with a final score of $k\alpha_1 - (k-1)\alpha_{1+i} - \alpha_{m-r(i)} + (k-1)\alpha_{1+i} + \alpha_{m-r(i)} = k\alpha_1$ as well.

For the converse, let $C \subseteq M$ be a set of at most k votes whose addition lets p win. If this is not a cover, then there is some s_i that is vetoed in none of the added votes. We now compare the points of p and s_i.

- p gains α_1 points in each of the $\|C\|$ additional votes, so p ends up with exactly $\|C\|\alpha_1$ points.

- Since s_i is not vetoed in any new vote, s_i gains α_{1+i} points in each added vote and thus ends up with $k\alpha_1 - (k-1)\alpha_{1+i} - \alpha_{m-r(i)} + \|C\|\alpha_{1+i}$ points.

Since $\|C\| \leq k$, $\alpha_1 \geq \alpha_{i+1}$ and $\alpha_{i+1} \geq \alpha_{3k+1} > \alpha_{m-2} \geq \alpha_{m-r(i)}$, it follows that $k\alpha_1 - (k-1)\alpha_{1+i} - \alpha_{m-r(i)} + \|C\|\alpha_{1+i} - \|C\|\alpha_1$
$= \underbrace{(k - \|C\|)(\alpha_1 - \alpha_{1+i})}_{\geq 0} + \underbrace{\alpha_{1+i} - \alpha_{m-r(i)}}_{>0} > 0$. So s_i beats p if C is not a cover; since by assumption adding C makes p win, C must be a cover. □

In a similar way, we can prove an analogous result for all scoring rules that "can implement" 4-approval in the sense that being voted in one of the first 4 positions is strictly better than being voted in most "later" positions. The proof of the following result is very similar to the proof of Theorem 4.8, except that an additional argument is needed to ensure that the favorite candidate cannot be made a winner with less than k additional voters.

Theorem 4.9. *Theorem 4.8 also holds when the condition $\alpha_{k+1}^{f,m} > \alpha_{m-2}^{f,m}$ is replaced with $\alpha_4^{f,m} > \alpha_{m-3k+1}^{f,m}$.*

As mentioned above, we now put the two reductions above together to obtain the NP-hardness result of this section, i.e., to prove that f-CCAV is NP-complete as soon as there is a number m where the coefficients of f satisfy $\alpha_4^{f,m} > \alpha_{n-2}^{f,m}$. If this condition is true, then we know that one of the inequalities $\alpha_4^{f,m} \geq \alpha_5^{f,m} \geq \cdots \geq \alpha_{m-3}^{f,m} \geq \alpha_{m-2}^{f,m}$ is in fact strict. Depending on the position of this strict inequality, we choose which reduction to apply: If the strict inequality appears "close" to the first candidate, then the first "few" positions are strictly better than "most," and the system can "simulate" k-approval for some $k \geq 4$. On the other hand, if the strict inequality appears "close" to the last candidate, then the last "few" positions are worse than "most," and we can similarly "simulate" k-veto for some $k \geq 3$.

Theorem 4.10. *f-CCAV is NP-complete for every* FP*-uniform pure $\mathbb{Q}$-generator f with $\alpha_4^{f,m} > \alpha_{m-2}^{f,m}$ for some m.*

4.2.3 "Few" Different Coefficients

We now study pure generators f *not* covered by Theorem 4.10, i.e., where $\alpha_4^{f,m} \leq \alpha_{m-2}^{f,m}$ for all m. Then for $m \geq 6$, $\alpha^{f,m}$ is of the form $(\alpha_1^{f,m}, \alpha_2^{f,m}, \alpha_3^{f,m}, \alpha_4^{f,m}, \ldots, \alpha_4^{f,m}, \alpha_5^{f,m}, \alpha_6^{f,m})$. The reductions above cannot work in this case, since there are no 3 positions "worse than most" and no 4 positions "better than most."

Due to Theorem 3.3, we can regard f equivalently as flexible $\mathbb{N}$-generators or as pure $\mathbb{Q}$-generators. For the latter representation, purity requires that all coefficients from $\alpha^{f,m}$ also appear in $\alpha^{f,m+1}$. So the above numbers $\alpha_1, \ldots, \alpha_6$ do not depend on m. We can use a fixed affine transformation for these finitely many coefficients and, using Proposition 2.1, rewrite all coefficients as natural numbers.

Our next hardness result concerns a generalization of 3-approval. Recall from Theorem 4.3 that CCAV for 3-approval itself, i.e., the generator $(\alpha, \alpha, \alpha, 0, \ldots, 0)$, is solvable in polynomial time. In Theorem 4.4, we proved a generalization of 2-approval to still give a polynomial-time solvable CCAV-problem. We now show that the analogous generalization of 3-approval leads to NP-completeness.

Theorem 4.11. *Let $\alpha \geq \beta \geq \gamma > 0$ and $\alpha \neq \gamma$. Let f be the generator giving $(\alpha, \beta, \gamma, 0, \ldots, 0)$. Then f-CCAV is NP-complete.*

Proof. Let M be the set from an instance of 3DM with $\|M\| = n$, and let $k = \|X\| = \|Y\| = \|Z\|$ (recall X, Y, and Z must be pairwise disjoint). We use the candidates $X \cup Y \cup Z \cup \{p\} \cup \{S_i, S_i' \mid S_i \in M\}$ and a dummy candidate d to be able to apply Lemma 4.7. We use the lemma to set up the registered votes such that the resulting relative scores are as follows: $\mathit{score}(p) = \alpha + 2\gamma$, $\mathit{score}(c) = (n+2k)\beta + 2\gamma$ for all $c \in X \cup Y \cup Z$, $\mathit{score}(S_i) = (n+2k)\beta + \min(\alpha, 2\gamma)$, and $\mathit{score}(S_i') = (n+2k)\beta + \alpha + \gamma$ for each $S_i \in M$. Further, $\mathit{score}(d) < -(n+2k)\alpha_1$. For each $S_i = (x, y, z)$, we introduce four unregistered voters voting as follows:

$x > p > S_i > \dots$.

$y > p > S_i > \dots$.

$z > p > S'_i > \dots$.

$S_i > p > S'_i > \dots$.

We show that p can be made a winner of the election by adding at most $n + 2k$ voters if and only if the 3DM-instance is positive, i.e., there is a set $C \subseteq M$ with $\|C\| = k$ and for $S_i \neq S_j \in C$, S_i and S_j differ in all three components.

First assume that there is such a cover. In this case, p can be made a winner of the election by adding the following voters: For each $S_i = (x, y, z) \in C$, we add the votes $x > p > S_i$, $y > p > S_i$, and $z > p > S'_i$. For each $S_i = (x, y, z) \notin I$, we add the vote $S_i > p > S'_i$. Note that this adds exactly $3k + (n - k) = n + 2k$ votes. Adding these votes results in the following scores:

- p gains β points in each added vote, so p gains $(n + 2k)\beta$ points and p's final score is $\alpha + 2\gamma + (n + 2k)\beta$,
- each candidate in $X \cup Y \cup Z$ gains α points, leading to a final score of $\alpha + \underbrace{(n + 2k)\beta + 2\gamma}_{\text{previous score}}$ as well,
- for each $S_i \in C$, we have that $score(S_i) = \underbrace{(n + 2k)\beta + \min(\alpha, 2\gamma)}_{\text{previous score}} + 2\gamma \leq (n + 2k)\beta + \alpha + 2\gamma$, which again is the score of p.
- for each $S_i \notin C$, we have $score(S_i) = (n + 2k)\beta + \min(\alpha, 2\gamma) + \alpha \leq (n + 2k)\beta + 2\gamma + \alpha$, equal to the score of p.
- for each S'_i (independent of whether $S'_i \in C$ or $S'_i \notin C$), we have $score(S'_i) = \underbrace{(n + 2k)\beta + \alpha + \gamma}_{\text{previous score}} + \gamma = (n + 2k)\beta + \alpha + 2\gamma$, again this is the score of p.

Thus all candidates tie and so in particular, p is a winner of the election.

For the converse, assume that p can be made a winner by adding at most $n + 2k$ voters. Since each S'_i initially beats p, at least one vote is added. Thus there is a candidate $c \in X \cup Y \cup Z$ with $score(c) \geq (n+2k)\beta+2\gamma+\alpha$, or some S'_i with $score(S'_i) \geq (n+2k)\beta+\alpha+2\gamma$. In both cases, we need to add at least $n+2k$ voters to ensure that p has at least $\alpha+2\gamma+(n+2k)\beta$ points as well.

Since $n + 2k$ votes are added, and each of these votes gives points to p and 2 other candidates, there are $2n + 4k$ positions awarding points in the added votes that are filled with (not necessarily different) candidates other than p. Each of the $3k$ candidates from $X \cup Y \cup Z$ can only gain α points without beating p in the election, so each of these can

fill at most one of these $2n + 4k$ positions. So at least $2n + k$ positions must be filled by (again, not necessarily different) candidates from $\{S_i, S'_i \mid 1 \leq i \leq n\}$. Each S'_i can appear at most once in the third position without beating p. Since there are n candidates of the form S'_i, it follows that there must be $n + k$ occurrences of candidates S_i in the first three positions of the added votes. Since no S_i can gain $\alpha + \gamma$ points without beating p,[2] each S_i can either appear in a vote $S_i > p > S'_i$, or in up to two votes of the form $c > p > S_i$ with $c \in X \cup Y$. ($S_i = (x, y, z)$ cannot appear in three of these, since then one of x and y would gain too many points.) So the only way to fill $n + k$ positions with candidates of the form S_i is having $2k$ occurrences of S_i in the third place, and $n - k$ occurrences of S_i in the first place. In order to fill all positions, each S'_i has to appear once in the final position, and due to the above, $n - k$ of these occurrences are in a vote of the form $S_i > p > S'_i$. Thus there are k votes of the form $z > p > S'_i$. It follows that there are $3k$ votes added that vote a candidate from $X \cup Y \cup Z$ in the first position, and $n - k$ voters are added that vote some S_i first. Since no S_i may appear both in first and in last position, and each S'_i may appear only once, and each x_i, y_i, and z_i may gain only α points, it follows that the added votes correspond to a cover. □

We also have proved the following cases NP-complete.

Theorem 4.12. *The problem f-CCAV is NP-complete if f is one of the following pure generators:*

1. $f = (\alpha_1, \alpha_2, \alpha_3, \alpha_4, \ldots, \alpha_4, \alpha_5, 0)$ *with* $\alpha_2 > \alpha_4 > 0$.
2. $f = (\alpha_1, \alpha_2, \ldots, \alpha_2, 0)$ *with* $\alpha_1 \notin \{\alpha_2, 2\alpha_2\}$, $\alpha_2 > 0$.
3. $f = (\alpha_1, \alpha_2, \ldots, \alpha_2, \alpha_5, 0)$ *with* $\alpha_1 > \alpha_2 > \alpha_5$.
4. $f = (\alpha_1, \ldots, \alpha_1, \alpha_5, 0)$ *with* $\alpha_1 > \alpha_5 > 0$.

4.3 Proof of Dichotomy Theorem

We now use the individual results from Sections 4.1 and 4.2 to prove our main dichotomy result, Theorem 4.2:

Proof. The polynomial cases follow from Theorems 4.3, 4.4, and 4.5, we prove hardness. If $\alpha_4^{f,m} > \alpha_{m-2}^{f,m}$ for some m, hardness follows from Theorem 4.10. So assume $\alpha_4^{f,m} = \cdots = \alpha_{m-2}^{f,m}$ for all $m \geq 6$. As argued in the discussion after Theorem 4.10, we assume $\alpha^{m,f} = (\alpha_1, \alpha_2, \alpha_3, \alpha_4, \ldots, \alpha_4, \alpha_5, \alpha_6)$ for each $m \geq 6$. Due to Proposition 2.1, we can assume $\alpha_6 = 0$. We reduce the number of relevant coefficients from 5 to 3:

- If $\alpha_4 = 0$, then, since f does not generate 3-approval and is not equivalent to f_4, $\alpha_1 > \alpha_3 > 0$. Hardness follows from Theorem 4.11.

[2] To see this, we compute the difference between the score of S_i after gaining $\alpha + \gamma$ points and that of p after gaining $(n+2k)\beta$ points. This value is $(n+2k)\beta + \min(\alpha, 2\gamma) + \alpha + \gamma - \alpha - \gamma - (n+2k)\beta = \min(\alpha, 2\gamma) - \gamma$. Since $\alpha > \gamma$ and $\gamma > 0$, this value is strictly positive, so S_i indeed beats p if S_i gains $\alpha + \gamma$ points.

- If $\alpha_2 > \alpha_4 > 0$, hardness follows from Theorem 4.12.1.

So assume $\alpha_2 = \alpha_3 = \alpha_4 > 0$, i.e., f is of the form $(\alpha_1, \alpha_2, \ldots, \alpha_2, \alpha_5, 0)$. We make a further case distinction:

- If $\alpha_2 = \alpha_5$, then since f does not generate 1-veto, we know that $\alpha_1 \neq \alpha_5 = \alpha_2$. Since f is not equivalent to $(2, 1, \ldots, 1, 0)$, we know that $\alpha_1 \neq 2\alpha_2$. Thus NP-hardness follows from Theorem 4.12.2.

- If $\alpha_2 > \alpha_5$, then depending on whether $\alpha_1 > \alpha_2 > \alpha_5$ or $\alpha_1 = \alpha_2 > \alpha_5$, hardness follows from Theorem 4.12.3 or Theorem 4.12.4 (note that in the latter case, we know that $\alpha_5 \neq 0$, since f does not generate 2-veto). □

Acknowledgments We thank the AAAI 2014 reviewers for helpful comments and suggestions.

References

J. Bartholdi, III, C. Tovey, and M. Trick. How hard is it to control an election? *Mathematical and Computer Modeling*, 16(8/9):27–40, 1992.

D. Baumeister and J. Rothe. Taking the final step to a full dichotomy of the possible winner problem in pure scoring rules. *Information Processing Letters*, 112(5):186–190, 2012.

N. Betzler and B. Dorn. Towards a dichotomy of finding possible winners in elections based on scoring rules. *Journal of Computer and System Sciences*, 76(8):812–836, 2010.

N. Betzler, R. Niedermeier, and G. Woeginger. Unweighted coalitional manipulation under the Borda rule is NP-hard. In *Proceedings of the 22st International Joint Conference on Artificial Intelligence*, pages 55–60. AAAI Press, July 2011.

J. Davies, G. Katsirelos, N. Narodytska, and T. Walsh. Complexity of and algorithms for Borda manipulation. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence*, pages 657–662. AAAI Press, August 2011.

P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. How hard is bribery in elections? *Journal of Artificial Intelligence Research*, 35:485–532, 2009.

P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. Weighted electoral control. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems*, pages 367–374. International Foundation for Autonomous Agents and Multiagent Systems, May 2013.

E. Hemaspaandra and L. Hemaspaandra. Dichotomy for voting systems. *Journal of Computer and System Sciences*, 73(1):73–83, 2007.

J. Hopcroft and J. Ullman. *Introduction to Automata Theory, Languages, and Computation.* Addison-Wesley, 1979.

A. Lin. *Solving Hard Problems in Election Systems.* PhD thesis, Rochester Institute of Technology, Rochester, NY, 2012.

W. Riker. *The Art of Political Manipulation.* Yale University Press, 1986.

A. Taylor. *Social Choice and the Mathematics of Manipulation.* Cambridge University Press, 2005.

Appendix

The appendix is structured as follows:

- In Section A, we provide additional discussion and full proofs for the results on descriptional richness and pure scoring rules contained in Sections 2 and 3 of this paper.
- Section B contains the proofs that were omitted from Section 4, i.e., the dichotomy result, and also provides additional discussion of "partial reductions."

A Omitted Proofs and Discussion from Sections 2 and 3

A.1 Proof of Proposition 2.1

Proposition 2.1. *Let $m \geq 1$ and let α and α' be m-position scoring vectors over $\mathcal{T} \subseteq \mathbb{Q}$. Then α and α' have the same winner sets on each m-candidate election if and only if α and α' both have the same normalized version.*

Proof. The "if" direction follows from Observation 2.2 of Hemaspaandra and Hemaspaandra (2007), as noted in Betzler and Dorn (2010).

Let us prove the "only if" direction. This result seems so natural and important that it feels as if it should be a folk theorem, although we don't know of it as such; but in any case, since the result is crucial to this paper, we include a construction that clearly and explicitly establishes this claim. Let D and D' be arbitrary, fixed length-m scoring vectors over N such that their normalizations, A and A', differ. We will construct an m-candidate election in which the winner sets under A and A' differ.

Since by the "if" direction of the present proposition A and D always yield the same winner set, and also by the "if" direction of the present proposition A' and D' always yield the same winner set, we may conclude that our constructed election has different winner sets under D and D'.

Let the components of the normalized scoring vector A be $(\alpha_1, \alpha_2, \ldots, 0)$, and let the components of the normalized scoring vector A' be $(\alpha'_1, \alpha'_2, \ldots, 0)$.

If $m = 1$, $A \neq A'$ is impossible. If $m = 2$, $A \neq A'$ exactly if one of them is $(0, 0)$ and one is $(1, 0)$, and these easily can be seen to give different winners on some inputs. Thus we from now on in this proof assume that $m \geq 3$.

A scoring vector is trivial if all its coefficients are the same, e.g., $(0, 0, 0)$. If A and A' are both trivial, then $A = A'$, and if exactly one is trivial then building an input separating them is easy. Thus we from now on in this proof assume that both A and A' are nontrivial.

Recall that both A and A' are normalized. We will in polynomial time "align" them, i.e., we will scale them so that the first components of A and A' become identical. In particular, to align them we will multiply A by α'_1 and we will multiply A' by α_1. So the first coefficient of each will now be $\alpha_1\alpha'_1$. Let us rename the thus-scaled vectors (each equivalent to its original vector as to what winner sets it gives, due to Observation 2.2 of Hemaspaandra and Hemaspaandra (2007)) as $B = (\beta_1, \beta_2, \ldots, \beta_{m-1}, 0)$ and $B' = (\beta'_1, \beta'_2, \ldots, \beta'_{m-1}, 0)$.

Let γ be the least i such that $\beta_i \neq \beta'_i$. W.l.o.g., assume $\beta_i > \beta'_i$. $\gamma = 1$ is impossible since $\beta_1 = \beta'_1 = \alpha_1\alpha'_1$. So $2 \leq i \leq m - 1$.

We will now specify an election in which B and B' have different winner sets.

Our candidates will be a, b, and "dummy" candidates $d_1, \ldots, d_{m-2}$. We will ensure that only a and b are serious contenders for winning.

Let s_1 be a shorthand for "$d_1 > d_2 > \cdots > d_{m-2}$," let s_2 be a shorthand for "$d_2 > d_3 > \cdots > d_{m-2} > d_1$," and so on, up to s_{m-2} being a shorthand for "$d_{m-2} > d_1 > \cdots > d_{m-3}$."

Our vote set is as follows:

1. For each i, $1 \leq i \leq m - 2$, we will have H votes "$a > b > s_i$" and we will have H votes "$b > a > s_i$." H's exact value will be specified later in the proof, but will be chosen to be so large that a and b are the only serious contenders.

2. β_γ votes of the form "a is the top choice and b is the last choice (and the other choices will be irrelevant to our proofs, but for specificity let us say they are filled in in lexicographical order)."

3. β_1 votes of the form "b is the γth choice and a is the last choice (and the other choices will be irrelevant to our proofs, but for specificity let us say they are filled in in lexicographical order)."

That ends our specification of the votes.

Let us tally up the points that each candidate gets under B and under B'. Clearly, and keeping in mind that $\beta_1 = \beta'_1$, we have: $score_B(a) = (m-2)H(\beta_1 + \beta_2) + \beta_1\beta_\gamma + 0$, $score_{B'}(a) = (m-2)H(\beta'_1 + \beta'_2) + \beta_\gamma\beta'_1 + 0$, $score_B(b) = (m-2)H(\beta_1 + \beta_2) + \beta_1\beta_\gamma + 0$, $score_{B'}(b) = (m-2)H(\beta'_1 + \beta'_2) + \beta_1\beta'_\gamma + 0$, $score_B(d_i) \leq (m-3)(2\beta_3) + \beta_\gamma\beta_2 + (\beta_1)^2$, and $score_{B'}(d_i) \leq (m-3)(2\beta'_3) + \beta_\gamma\beta'_2 + \beta_1\beta'_1$. We now set the value of H, namely to be $H = 2\beta_1$. Keeping in mind that $\beta_1 = \beta'_1 \geq 1$, it is easy to see that this choice of H ensures that for each i, $1 \leq i \leq m - 2$, $\min(score_B(a), score_B(b)) > score_B(d_i)$ and $\min(score_{B'}(a), score_{B'}(b)) > score_{B'}(d_i)$. So we have ensured that, under B and under B', a and b have more points than any d_i.

Under B', note that a is the one and only winner (recall $\beta_1 = \beta'_1$ and $\beta_\gamma > \beta'_\gamma$). But our votes ensure that under B, a and b tie as the (sole) winners.

So the votes we gave show that B and B' (and thus D and D') have different winner sets on some example, namely, the above example. □

A.2 Proof of Theorem 3.1

Theorem 3.1. *If $T_2(m)$ is a fully time-constructible function and $\limsup_{n\to\infty} \frac{T_1(m)\log T_1(m)}{T_2(m)} = 0$, then there is an election rule in* **FDTIME[$T_2(m)$]-*uniform*-{0, 1}-PSR** *that is not in* **FDTIME[$T_1(m)$]-*uniform*-$\mathbb{Q}$-GSR**.

Proof. Rather than proving Theorem 3.1, we will prove a slightly weaker result, "Theorem X," and then will explain how to modify that proof to establish Theorem 3.1.

> **Theorem X** There is an integer constant $k > 0$ such that if $T_2(m)$ is a fully time-constructible function, and $\limsup_{n\to\infty} \frac{T_1(m)\log T_1(m)}{T_2(m)} = 0$, and $(\forall m)[T_2(m) \geq k(m+1)]$, then there is an election rule in **FDTIME[$T_2(m)$]-uniform-{0, 1}-PSR** that is not in **FDTIME[$T_1(m)$]-uniform-Z-GSR**.

We will prove Theorem X by describing how to appropriately adjust the classic presentation by Hopcroft and Ullman (1979), henceforward "HU," of the proof of the deterministic time hierarchy theorem (their Theorem 12.9); that proof of HU can be found as pages 297–298 of that book, and that proof itself draws also on the framework of that book's earlier proof of the deterministic space hierarchy theorem. We will assume that the reader is familiar with those deep, classic proof presentations; anyone not expert in complexity will want to either first master that proof framework or just skip the present proof.

So, suppose we are given T_1 and T_2 satisfying the conditions of Theorem X. We'll later specify k, but k must not and will not depend on T_1 or T_2.

Our goal is to show that there is an FDTIME[$T_2(m)$]-uniform-$\{0,1\}$-PSR generator, call it f, whose rule's winner set is not obtained by *any* FDTIME[$T_1(m)$]-uniform-Z-GSR.

Generally, we just follow the overall architecture of the HU proof, but the differences are as follows. First, we will be computing and diagonalizing against functions. So our enumeration of machines will be an enumeration of all function-computing Turing machines (i.e., machines with an output tape such that when they halt, whatever is on the output tape is viewed as being the output). Like HU, we will assume that our enumeration of such machines has the property that if w encodes a machine, then $1^k w$ encodes exactly the same machine; "padding" by adding 1^* as a prefix to a machine's coding does not change the machine encoded. (Unlike HU, who truly need this due to their making a liminf claim, we merely need each machine to appear infinitely often in the enumeration, since for other reasons we are simply making a limsup claim. However, this padding property certainly is a fine way of achieving what we need.)

Another difference is that we need to be a pure scoring rule. So each length rule must appropriately link to and extend the rule from the previous length. As mentioned in the

text immediately after Theorem 3.1, doing so in the obvious fashion would seem to add a multiplicative factor of m, but instead we use the trick mentioned there to avoid this. That is, for each odd natural number m, $f(0^m)$ will be $1^{\lfloor m/2 \rfloor}0^{\lfloor m/2 \rfloor+1}$ (e.g., (0), $(1,0,0)$, $(1,1,0,0,0)$, etc.). This is purely mechanical, and does not involve any diagonalizing. (One might worry that for some very small values of m we might not even have time to realize that m was odd and write the appropriate output; that worry is because one might worry that for some small m we may not have time greater than $m+1$ and that is not enough time to both see what m is, and that it is odd, and to already just before halfway through it have known that we are just before halfway through it and to have switched what we are outputting from 1s to 0s. However, the $(\forall m)[T_2(m) \geq k(m+1)]$ assumption of Theorem X will ensures us that we do have that time.)

Now, at each even length m, we'll try to do a diagonalization, if we have time. Our framework is that we will always have as our vector at this length either $1^{\lfloor m/2 \rfloor+1}0^{\lfloor m/2 \rfloor+1}$ or $1^{\lfloor m/2 \rfloor}0^{\lfloor m/2 \rfloor+2}$ (e.g., at length 4, either 1100 or 1000; note that in this proof we will quietly go back and forth notationally between, for example, $(1,1,0,0)$ and 1100 and 1^20^2). Basically, we want to ensure that if whatever FDTIME$[T_1(m)]$-uniform-Z-GSR generator we ("we" are the function f) at stage m are trying to diagonalize against outputs a vector whose winner set is the same as that given by $1^{\lfloor m/2 \rfloor+1}0^{\lfloor m/2 \rfloor+1}$ then we will output the vector $1^{\lfloor m/2 \rfloor}0^{\lfloor m/2 \rfloor+2}$ and otherwise we will output the vector $1^{\lfloor m/2 \rfloor+1}0^{\lfloor m/2 \rfloor+1}$. To do this, we immediately write onto our output tape the vector $1^{\lfloor m/2 \rfloor+1}0^{\lfloor m/2 \rfloor+1}$. We do so so that if we run out of time in our diagonalization on this input, we at least have one of the two legal vectors that keep our rule pure within our framework. And we do have time to write this vector, thanks to the $(\forall m)[T_2(m) \geq k(m+1)]$ assumption of Theorem X.

Now, as to the diagonalization, it goes as follows. If $m \geq 1$ is even, then strip away the leading 1 of m in binary and the trailing 0 of m in binary, and call the string that remains w. (We strip the trailing 0 because we diagonalize only at even m, but we want the entire set of such strings created to exactly equal $\{0,1\}^*$.) We will view w as the encoding of a Turing machine from our enumeration of function-computing TMs. We'll run it (unless we run out of time: as per the entire HU framework, we'll be using a separate tape and the fully time-constructible nature of T_2 to enforce a time cutoff; actually, we'll use two separate tapes, since our time cutoff is $\max(T_2(m), k(m+1))$ on input 0^m to get the length-m vector it outputs, call it v (if it halts and has the wrong length output, then it clearly isn't even a valid GSR). We must then efficiently evaluate whether that vector has the same winner set as would $1^{\lfloor m/2 \rfloor+1}0^{\lfloor m/2 \rfloor+1}$. One might think to do so we would have to normalize v, which includes gcd's and other time-eating computations, but that is not so. To tell if v has the same winner set as $1^{\lfloor m/2 \rfloor+1}0^{\lfloor m/2 \rfloor+1}$, we need only test whether v is of the form $a^{\lfloor m/2 \rfloor+1}b^{\lfloor m/2 \rfloor+1}$, where a and b are members of Z and $a > b$. If so they have the same winner set, and if not they don't. Ignoring for the moment the cost of getting v, this test is an easy linear-time test on a multihead, multitape TM. If we find that v does have the same winner set as $1^{\lfloor m/2 \rfloor+1}0^{\lfloor m/2 \rfloor+1}$, then we on our output tape just overwrite the $\lfloor m/2 \rfloor$th character, changing it from a 1 to a 0 (for example, if $m = 4$, we'd overwrite the second bit to change our placeholder 1100 into 1000). So, if we have time for our simulation of w to

complete and to evaluate whether the v has the same winner set as $1^{\lfloor m/2 \rfloor+1}0^{\lfloor m/2 \rfloor+1}$, and if it is, have time to overwrite the one bit, then we have successfully diagonalized against the machine w.

On the other hand, what if we run out of time? No problem. If the machine associated with w happens to run in time $T_1(m)$ for all m (and we're doing all w, so some will have that time behavior and some won't), then—since w will appear infinitely often again in our construction, as it will also appear as $1w$, as $11w$, and so on—due to the limsup assumption we eventually *will* have the time to fully run the diagonalization (the simulation and our on-the-cheap comparison of winner sets and our bit-fixing). The reason is that the log multiplicative overhead in the limsup is (see HU) enough to do the machine simulation, and the limsup ensures that for any constant c our T_2 for every m beyond some point will satisfy $T_2(m) > cT_1(m)\log T_1(m)$.[3]

Let us discuss the constant k. k is simply there because for small values of m, we may not have enough time to see the input, and write the appropriate starting string to our output tape. (All we know is we have at least $m+1$ steps. That isn't enough, as out input is of the form 0^m so until we hit the right-end marker, all we see is 0s, but the starting string we need to put on the output tape as a placeholder is roughly half 1's and half 0's, so if we just get $m+1$ steps, we don't know roughly halfway through the 0's that we are roughly halfway through and that we have to switch from outputting 1's to outputting 0's. The fact that asymptotically we have at least $T_1(m)\log T_1(m)$ time, and thus at least $m\log m$ time, doesn't help us for small values of m, since if we run out of placeholder time at even one value of m we've violated purity of our alleged PSR.) However, we can set a value of k that allows us to do the placeholder writing (and on any natural TM framework, it will be a very small constant k; go to the end of the string 0^m, see if it is even, and then output the right placeholder string, for example by having on an extra dummy tape basically build a string of just under $m/2$ 1's that we can use to trigger our switch-over point as we write the placeholder output), for even small values. Note that k depends just on the TM framework, not on T_1 or T_2.

All that remains is how to move from our proof of Theorem X to a proof of Theorem 3.1.

[3]Readers familiar with the time hierarchy theorem may wonder why we here use a limsup but the HU theorem uses a liminf. The reason is subtle. Briefly put, in HU, once a machine w is being diagonalized against at some length, it will be (in its sibling forms $1w$, $11w$, etc.) diagonalized against at *every* greater length. So a liminf suffices, as liminf ensures that at *some* (indeed, infinitely many, though they don't really need that) greater lengths we get enough time overhead to diagonalize, and that is all HU needs. (Warning: The HU proof itself sets up the right machinery and construction for liminf to suffice, but then inside its proofs, seems to forget this and by such phrases as "infinitely often" on its p. 298 line 8 and "has arbitrarily long encodings" on page 298 and 299, say things that would require limsup to support. However, the HU construction supports liminf, and one can change the two mentioned phrases to "almost everywhere" and "at every length onward once it first occurs" and the thus adjusted proof becomes correct.) In contrast, PSRs take as input 0^m, and so at each length are diagonalizing against just *one* Turing machine. And so we do know that once we face off against w (say, on input $0^{(1w)_{binary}}$), we'll infinitely often see its identical siblings at longer lengths (such as on input $0^{(11w)_{binary}}$, but note that that is about twice as long, not one longer), but we don't know we'll have them at every length. Basically, since our input is 0^m, we get to attempt just one diagonalization at each length. Happily, under a limsup assumption, we're safe and fine, since at every sufficiently long length, we have the headroom to diagonalize.

The previous paragraph's contortions will actually already hint to the reader the path to doing this. Namely, suppose we now must operate without the crutch of having $(\forall m)[T_2(m) \geq k(m+1)]$. Our salvation is the limsup claim, which ensures that eventually, T_2 becomes very big. So all we need to do is find a way to mark time until that kicks in, and by marking time, we have to make do with $m+1$ steps as that is all we know we always have. We can't such marking time within the framework of the particular PSR we've been using above, because as noted above, one would seem to have to be able to know the midpoint of 0^m in real time as one was passing it, and that is impossible. The workaround is as follows. We change the PSR that we'll use. In particular, for each T_1 and T_2, there exists some odd value m' beyond which the limsup ensures that we do have easily enough time headroom. So what we'll do is our PSR will simply be 0^m at each length $m \leq m'$ (note that we can do this in time $m+1$, at length m), and from then on, it will mimic our trick except with those extra 0's, namely, we will fixing our vector at each odd length, $m > m'$, to be $1^{\lfloor (m-m')/2 \rfloor} 0^{\lfloor (m-m')/2 \rfloor + 1 + m'}$, and at even lengths $m > m'$, our vector will be either $1^{\lfloor (m-m')/2 \rfloor + 1} 0^{\lfloor (m-m')/2 \rfloor + 1 + m'}$ or $1^{\lfloor (m-m')/2 \rfloor} 0^{\lfloor (m-m')/2 \rfloor + 2 + m'}$. This allows us to remove the requirement that $(\forall m)[T_2(m) \geq k(m+1)]$, and so we have established Theorem 3.1. Note that m' will depend on T_1 and T_2, but that is perfectly legal; it does not undermine the proof. □

We finish with some brief, technical comments about possible extensions and alternate proofs. First, our diagonalization is carried out the world that it is in, namely, the world of functions. We mention in passing that since the PSR we build is variable only as to which of two possible vectors it has at each even length, the information behind its choice is in effect a set—indeed, a set of the form $A \subseteq (11)^*$, and so a particular type of tally set. One could use this to make the proof more set-focused. But one would still have to account for overhead time and so on, and we feel it is more natural to simply have the proof be in the world the items in question are inhabiting. Second, we mention that it is possible that one could turn Theorem 3.1's limsup into a liminf, thus making the theorem stronger. However, for most natural time functions T_1 and T_2, the liminf and the limsup are equal, and so we do not consider this an important direction, especially as it would likely make the proof far more complex if attempted. We do mention that the natural path to use to attempt such an improvement would be, once one first was trying to diagonalize against a given machine, to attempt diagonalizing against that machine at every length from then on until one had successfully diagonalized against it. But we warn that to even know what one was currently trying to diagonalize against would require recomputing one's history, which itself takes time and can interfere with the proof. Worse, our current setup doesn't do any diagonalizations at even lengths, but the liminf could go to zero just due to even lengths, and that would not at all help us *ever* have time to diagonalize; so our entire zigzag framework becomes poisonous, yet abandoning it loses its advantage (already lost anyway in the type of new construction we are speculating about) of avoiding history rebuilding.

A.3 Proof of Theorem 3.2

In the following theorem, nonuniform simply means that the corresponding generators are not required to be computable in any specific complexity class. In fact, the generators may be uncomputable.

Theorem 3.2. $\mathbb{Q}_{\geq 0}$**-PSR** $\not\subseteq$ ***nonuniform-*$\mathbb{Z}$**-PSR**, $\mathbb{Z}$**-PSR** $\not\subseteq$ ***nonuniform-*$\mathbb{Q}_{\geq 0}$**-PSR**, *and* $\mathbb{Q}$**-PSR** $\not\subseteq$ ***nonuniform-*$\mathbb{Q}_{\geq 0}$**-PSR** $\cup$ ***nonuniform-*$\mathbb{Z}$**-PSR**.

Proof. Let us first show that $\mathbb{Q}_{\geq 0}$**-PSR** $\not\subseteq$ **nonuniform-**$\mathbb{Z}$**-PSR**. Consider the following FP-uniform $\mathbb{Q}_{\geq 0}$-PSR generator f.
$f(0^1) = (1)$,
$f(0^2) = (\frac{3}{2}, 1)$,
$f(0^3) = (\frac{3}{2}, 1, \frac{1}{2})$,
$f(0^4) = (\frac{7}{4}, \frac{3}{2}, 1, \frac{1}{2})$,
$f(0^5) = (\frac{7}{4}, \frac{3}{2}, 1, \frac{1}{2}, \frac{1}{4})$,
$f(0^6) = (\frac{15}{8}, \frac{7}{4}, \frac{3}{2}, 1, \frac{1}{2}, \frac{1}{4})$,
$f(0^7) = (\frac{15}{8}, \frac{7}{4}, \frac{3}{2}, 1, \frac{1}{2}, \frac{1}{4}, \frac{1}{8})$,
etc.

Consider a nonuniform $\mathbb{Z}$-PSR type generator g that claims to have the same winner set as this on all instances. By the (clear) extension of Proposition 2.1 mentioned at the end of the paragraph that follows that result, the 5-candidate vector of g must normalize to $(6, 5, 3, 1, 0)$. The ordered gap pattern in this is: 1,2,2,1. So the actual 5-candidate vector (over Z) of g satisfies $2(\alpha_1^5 - \alpha_2^5) = (\alpha_2^5 - \alpha_3^5) = (\alpha_3^5 - \alpha_4^5) = 2(\alpha_4^5 - \alpha_5^5)$. Similarly, the 7-candidate vector of g must normalize to a vector having gap pattern: 1,2,4,4,2,1. Since that vector retains all 5 coefficients from the 5-candidate vector, with two additional coefficients appropriately added, by inspection of the possibilities, it is clear that the only possible extension that achieves this is one that adds exactly the coefficients $\alpha_1^5 + \frac{1}{2}(\alpha_1^5 - \alpha_2^5)$ and $\alpha_5^5 - \frac{1}{2}(\alpha_1^5 - \alpha_2^5)$ (and due to the 6-candidate vector, they must be added in the order just stated). That is, the vector extends to the right and left, but an amount half the "base" gap. However, note that if $\alpha_1^5 - \alpha_2^5$ is not a multiple of 2, this extension is already not legal over $\mathbb{Z}$, as it would have coefficients not in $\mathbb{Z}$. If $\alpha_1^5 - \alpha_2^5$ is a multiple of 2, then we indeed do have this valid extension of the 5-candidate vector. But by the same argument, the only possible 9-candidate vector again will have to be formed from the 7-candidate vector by extending to the right and left ends by half the then-current "base" gap. Note that that will not be possible unless $\alpha_1^5 - \alpha_2^5$ is a multiple of $2^2 = 4$. And so on. Since there is some power of 2 that is the maximum power of two that divides $\alpha_1^5 - \alpha_2^5$, this process will eventual end in failure, i.e., there will be no valid vector over Z that properly extends the pattern g has trapped itself into. Thus, contrary to the claim made for g, g in fact cannot match the winner sets of f.

Let us now show that $\mathbb{Z}$**-PSR** $\not\subseteq$ **nonuniform-**$\boldsymbol{Q'}$**-PSR**. We will give the argument for this case more compactly, since the reader by now will be familiar with the flavor of arguments of this sort, from the above case. Consider the following FP-uniform $\mathbb{Z}$-PSR generator f.

$f(0^1) = (0)$,
$f(0^2) = (1, 0)$,
$f(0^3) = (1, 0, -1)$,
$f(0^4) = (2, 1, 0, -1)$,
$f(0^5) = (2, 1, 0, -1, -2)$,
$f(0^6) = (4, 2, 1, 0, -1, -2)$,
$f(0^7) = (4, 2, 1, 0, -1, -2, -4)$,
etc.

Consider a nonuniform $\mathbb{Q}_{\geq 0}$-PSR type generator g that claims to have the same winner set as this on all instances. Note that the gap pattern, at each length starting at 5, is that we have four identical, adjacent smallest-size (among the gaps that exist) gaps in the center, and then surrounding that the gaps grow by repeatedly doubling (with the extra one being on the "upper" side on even-length cases). Note that no single internal insertion anywhere maintains a four, identical, adjacent, smallest-size (among the gaps that exist) gap pattern. So we can only expand by going doubly-far to the right and left in each pair of extensions. But that means the values α_{2m+1}^{2m+1} will be decreasing at an exponentially increasing rate. And so since α_1^1 was finite, at some m the value α_{2m+1}^{2m+1} will necessarily have to be less than zero, and so will not be an element of $\mathbb{Q}_{\geq 0}$.

We'll sketch $\mathbb{Q}$**-PSR** $\not\subseteq$ **nonuniform-**$\mathbb{Q}_{\geq \mathbf{0}}$**-PSR** $\cup$ **nonuniform-**$\mathbb{Z}$**-PSR** even more briefly. It does not automatically follow from the earlier parts; claiming that would be like claiming that since not all integers are even and not all integers are odd, it follows that not all integers are odd integers or even integers. Nonetheless, by building a *construction* that uses both of the weaknesses exploited by the constructions of the preceding two parts, we can easily establish that $\mathbb{Q}$**-PSR** $\not\subseteq$ **nonuniform-**$\mathbb{Q}_{\geq \mathbf{0}}$**-PSR** $\cup$ **nonuniform-**$\mathbb{Z}$**-PSR**. In particular, it is easy to prove that the following FP-uniform $\mathbb{Q}$-PSR has the desired property. After setting down 4 equally spaced gaps, it in term makes re-doublingly large gaps going in the negative direction and re-halvingly large gaps going in the positive direction.
$f(0^1) = (2)$,
$f(0^2) = (2, 1)$,
$f(0^3) = (2, 1, 0)$,
$f(0^4) = (2, 1, 0, -1)$,
$f(0^5) = (2, 1, 0, -1, -2)$,
$f(0^6) = (2, 1, 0, -1, -2, -4)$,
$f(0^7) = (2\frac{1}{2}, 2, 1, 0, -1, -2, -4)$,
$f(0^8) = (2\frac{1}{2}, 2, 1, 0, -1, -2, -4, -8)$,
$f(0^9) = (2\frac{3}{4}, 2\frac{1}{2}, 2, 1, 0, -1, -2, -4, -8)$,
$f(0^{10}) = (2\frac{3}{4}, 2\frac{1}{2}, 2, 1, 0, -1, -2, -4, -8, -16)$,
$f(0^{11}) = (2\frac{7}{8}, 2\frac{3}{4}, 2\frac{1}{2}, 2, 1, 0, -1, -2, -4, -8)$,
etc.

By using the natural variants for this of both types of arguments used in the earlier parts of this proof, we can argue that this function can be accepted by neither a $\mathbb{Q}_{\geq 0}$-type generator nor a $\mathbb{Z}$-type generator. □

A.4 Proof of Theorem 3.3

Theorem 3.3. $\mathbb{N}$**-FPSR** $= \mathbb{Q}$**-FPSR** $= \mathbb{Q}_{\geq \mathbf{0}}$**-FPSR** $= \mathbb{Z}$**-FPSR** $= \mathbb{Q}$**-PSR**.

Proof. Theorem 3.3 is immediately clear, in light of Proposition 2.1 and the end of the paragraph following it, since our normalizations shifted each of these types to equivalent vectors over $\mathbb{N}$, which is the most restrictive of all these types.

The equality with $\mathbb{Q}$-PSR can easily be seen by inductively constructing, for every flexible generator f, an equivalent pure generator over the rationals. □

A.5 Proof of Theorem 3.4

Theorem 3.4. *There is an* FP*-uniform* PSR *that is not generated even by any nonuniform* $\text{PSR}_{\textit{norm-0}}$ *generator. On the other hand, every* FP*-uniform (respectively, nonuniform)* PSR *is generated by a* FP*-uniform (respectively, nonuniform)* PSR *generator that for all but a finite number of m has the property that the last coefficient, $\alpha_{m,m}$, is zero and the gcd of the nonzero coefficients in the length-m vector is one.*

Proof. Let us prove the first part of the theorem. Consider this FP-uniform $\mathbb{N}$-PSR, which we mention in passing is even an FP-uniform $\mathbb{N}$-$\text{PSR}_{\text{norm-gcd}}$. $f(0^1) = (5)$, $f(0^2) = (6, 5)$, $f(0^3) = (7, 6, 5)$, $f(0^4) = (8, 7, 6, 5)$, and $f(0^5) = (8, 7, 6, 5, 0)$. We won't use longer lengths here, but to make clear that this is a pure scoring rule, let us say that the remaining lengths are $f(0^6) = (8, 7, 6, 5, 0, 0)$, $f(0^7) = (8, 7, 6, 5, 0, 0, 0)$, and so on.

Consider an $\mathbb{N}$-$\text{PSR}_{\text{norm-0}}$ type generator g' that claims to have the same winner set as this on all instances. So by Proposition 2.1, it is clear that the 4-candidate vector of g' must be of the form $(3k, 2k, k, 0)$, for some $k \in \mathbb{N} - \{0\}$.

But clearly addition of one coefficient to this vector cannot yield anything that normalizes to $(8, 7, 6, 5, 0)$; one can see this easily by examining each of the legal insertion places and seeing that it cannot suffice. So the rule of g' differs from the original rule on some 5-candidate examples, by Proposition 2.1.

(We mention in passing that the exact same f we just used can also be easily argued to be accepted by no PSR generator that has the property that at each length, the gcd of its nonzero coefficients (if any) is one.)

Let us turn to the second claim of the theorem, namely, that every $\mathbb{N}$-PSR is generated by an $\mathbb{N}$-PSR generator that on all but at most a finite number of lengths m has the property that the last coefficient is zero and the gcd of the nonzero coefficients in the length-m vector is one.

Given $\mathcal{E}_1$, an FP-uniform N-PSR, fix an FP generator for that rule. Let $\mathit{Pool} = \{\alpha_j^m \mid m \geq 1 \wedge 1 \leq j \leq m\}$, i.e., it is the set of all coefficients.

If $\|\mathit{Pool}\| = 1$ we are done easily using the scoring vector sequence (0), $(0, 0)$, $(0, 0, 0)$, ...

So henceforth we assume that $\|\mathit{Pool}\| \geq 2$. Let $\min(\mathit{Pool}) = k$. Now, alter the generator from now on (and thus each α_j^i, and so indirectly we also have altered Pool) to subtract k

from each original α_j^i. So it is clear that, after the alteration, at all but a finite number of values of m we have $\alpha_m^m = 0$.

Let $j = \gcd(Pool)$, for the (altered) $Pool$.

Let us keep in mind that with PSRs, coefficients that occur at a length must also occur at each greater length. If $j = 1$, then there is some finite length at which the set of coefficients at that length has a gcd of one, and so at all greater lengths the coefficient set also has a gcd of one. If $j \geq 2$, then further re-alter our generator to, after subtracting k, divide by j. Applying the $j = 1$ argument to this, we again have that the gcd will be one at all but a finite number of lengths, and as dividing by j does not change the set of lengths at which $\alpha_m^m = 0$, we also have that for our re-altered generator α_m^m is zero for all but a finite number of m's. The re-altered generator is clearly polynomial-time computable if the original one is, so we have completed the proof of the theorem's second part for the FP-uniform case, and the result clearly also holds for the nonuniform case, by the same argument with the discussion of runtimes removed. □

B Omitted Proofs and Discussion from Section 4

B.1 Missing Pieces of the Proof of Theorem 4.4

Here we present the proofs of the two facts used in the proof of Theorem 4.4. We start with Fact 1:

Fact 1. *For $i \in \{1, 2\}$, given a set $S \subseteq V_1 \cup V_2$, it can be checked in polynomial time whether S can be extended, by adding at most $k - \|S\|$ voters from V_i to make p win.*

Proof. For each candidate $c \in C$, let s_c be the score of that candidate after the voters in V and the voters in R have been counted. For the remainder of the proof of Fact 1, we assume that $i = 1$, i.e., we are adding voters with p in the first position. Since the proof of Fact 1 does not use the fact that $\alpha \geq \beta$, the proof works in the same way for $i = 2$, i.e., adding voters voting p in the second position.

We can add at most $r = k - \|S\|$ voters from V_1. Clearly, k is bound by the size of the instance. Thus we can brute-force over every $j \in \{0, \ldots, r\}$ and check whether S can be extended with exactly j voters from V_1. To do this for a given j, we proceed as follows:

- First, we compute the number of points that p will have after adding j voters from V_1, this is $s_p + j\alpha$. We denote this number with s_p^{final}.
- For each candidate $c \in C$, we compute the maximal number m_c of votes giving points to c that can be added. This is the maximal number m_c with $s_c + m_c\beta \leq s_p^{final}$. If m_c is negative for some c, then S cannot be extended with exactly j voters from V_1.
- For each candidate $c \in C$, we add voters voting c in the second place while still possible, i.e., while all of the following hold:
  - the number of added voters up to now does not exceed j,

– the number of added voters voting c second does not exceed m_c,

– there are still voters available voting c second.

- If the resulting set is a solution to the control problem, accept. Otherwise, S cannot be extended with exactly j voters from V_1. □

The proof for Fact 2 follows:

Fact 2. *To make p win, it is never better to add ℓ voters from V_2 than adding ℓ pairwise different voters from V_1.*

Proof. Let S_1 (S_2) be the corresponding subset of V_1 (V_2). Let $c_1, \ldots, c_\ell$ be the candidates voted in the second place by the voters in S_1. We compare the effect that adding S_1 or S_2 has on the relationship between p and the c_i (clearly, for the relationship between p and candidates $c \notin \{c_1, \ldots, c_\ell\}$, adding S_1 is always better).

- When adding S_1, p gains $(\ell - 1)\alpha + (\alpha - \beta) = \ell\alpha - \beta$ points against each c_i.
- When adding S_2, p gains at most $\ell\beta$ points against each of the c_i.

Since $\ell(\alpha - \beta) \geq \beta$, we have that $\ell\alpha - \beta \geq \ell\beta$, and so adding S_1 is at least as good as adding S_2. □

B.2 Proof of Lemma 4.7

The proof Lemma 4.7 makes extensive use of the following construction:

Lemma B.1. *Given a scoring vector $(\alpha_1, \ldots, \alpha_m)$, and numbers $i \neq j$, $k \neq l \in \{1, \ldots, m\}$, there is a set V of m votes such that when counting the votes in V, the scores of all candidates is $A := \sum_{m=1}^{n} \alpha_m$, except:*

- *candidate c_i has $A + \alpha_k - \alpha_l$ points,*
- *candidate c_j has $A + \alpha_l - \alpha_k$ points.*

Proof. Without loss of generality, assume $i \leq j$ and $k \leq l$, otherwise we swap.

Let v_1 be a vote having c_i in position l and c_j in position k, the remaining candidates are ordered arbitrarily. For any $\ell \geq 1$, let $v_{\ell+1}$ be the vote obtained from v_ℓ by moving each candidate one position to the right, i.e., the candidate at position t in the vote v_ℓ ends at position $t + 1$ in the vote $v_{\ell+1}$ if $t + 1 \leq m$, and at position 1 if $t + 1 = m$. Now let V_0 contain the votes $v_1, \ldots, v_m$. Clearly, each candidate gains exactly A points in the votes V_0. Now, let V be obtained from V_0 by exchanging, in the vote v_0, the positions of candidate c_i and c_j. Then, relatively to the votes V_0, the points of c_i and c_j change as follows:

- c_i loses α_l points and gains α_k points,
- c_j loses α_k points and gains α_l points.

So the relative score is as required. □

The actual proof of Lemma 4.7 follows:

Lemma 4.7. *Given a scoring vector $(\alpha_1, \ldots, \alpha_m)$, and for each $c \in \{1, \ldots, m-1\}$, numbers $a_1^c, \ldots, a_m^c$ in signed unary, and a number k in unary, we can compute, in polynomial time, votes such that the scores of the candidates when evaluating these votes according to the scoring vector $(\alpha_1, \ldots, \alpha_m)$ are as follows: There is some o such that for each $c \in \{1, \ldots, m-1\}$, $score(c) = o + \sum_{i=1}^{m} a_i^c \alpha_i$, and $score(c) > score(m) + k\alpha_1$.*

Proof. The algorithm produces the set V by consecutive applications of Lemma B.1. We say that an application of the Lemma *transfers* $(\alpha_k - \alpha_l)$ points from c_j to c_i, since the relative score (in relation to all candidates except c_i and c_j) of c_i increases by $\alpha_k - \alpha_l$, while the score of c_j decreases by the same amount. Note that in order for such a "transfer" to be possible, neither candidate is required to actually have a nonzero score. We start with the empty set V, and then, for each $a_i^c \neq 0$ for some $c \in \{1, \ldots, m-1\}$, we proceed as follows:

- if $a_i^c > 0$, we apply Lemma B.1 a_i^c times to transfer α_i points from the candidate m to the candidate c,
- if $a_i^c < 0$, we apply Lemma B.1 $(m-2)a_i^c$ times to transfer α_i points from the candidate m to each candidate $c' \neq \{c, m\}$.

In the above steps, clearly the score of m is the lowest among all candidates. So to ensure that $score(i) > score(m) + k\alpha_1$, it suffices to transfer $(k+1)\alpha_1$ points from the candidate m to each candidate $c \in \{1, \ldots, m-1\}$. Clearly, the resulting scores are as required. □

B.3 Proof of Theorem 4.9 and Discussion of "Partial Reductions"

In the discussion surrounding Theorems 4.8 and 4.10, we described Theorem 4.8 and its "dual"—which we now state formallyas Theorem B.2—as "partial reductions." To clarify this unusual terminology, we now explain why we use this term. For a generator f that is covered in Theorem 4.10, it is not necessarily the case that hardness follows from Theorem 4.8 alone (or from Theorem B.2 alone). Rather, for each value of m resulting from a 3DM-instance, the reduction establishing NP-hardness (which is given in the proof of Theorem 4.10) needs to check which of the two cases (as to meeting the differing "$\alpha_{\ldots} < \alpha_{\ldots}$" conditions listed in Theorem 4.8 and Theorem B.2) is satisfied. Depending on which condition applies, either the reduction g from Theorem 4.8 or the reduction g from Theorem B.2 is used for this particular 3DM-instance. In this sense, Theorem 4.8 (and also its "dual," Theorem B.2) are "partial" reductions: For a given f, each of these theorems may only be "half" of the reduction needed to prove the NP-hardness of f-CCAV. Theorems 4.8 and B.2 together, with a simple additional observation, then give Theorem 4.10, which is the "complete" reduction to obtain NP-completeness of f-CCAV.

Let us turn from the above discussion to the formal statement of Theorem 4.9 and its proof: We recall the statement of the Theorem:

Theorem 4.9. *Theorem 4.8 also holds when the condition $\alpha_{k+1}^{f,m} > \alpha_{m-2}^{f,m}$ is replaced with $\alpha_4^{f,m} > \alpha_{m-3k+1}^{f,m}$.*

The complete statement of Theorem 4.9 is as follows. Note that again, we do not require the generator f to respect any purity conditions.

Theorem B.2. *Let f be an* FP*-uniform $\mathbb{Q}$-generator. Then there exists an* FP*-computable function g such that*

- *g takes as input an instance I_{3DM} of 3DM where $k = \|X\| = \|Y\| = \|Z\|$,*
- *g produces an instance I_{CCAV} of f-CCAV,*
- *If for $m = 6k$ we have that $\alpha_4^{f,m} > \alpha_{m-3k+1}^{f,m}$, then: I_{3DM} is a positive instance of 3DM if and only if I_{CCAV} is a positive instance of f-CCAV.*

Proof. The proof is very similar to the proof of Theorem 4.8 above, in particular we use the same set of candidates. Again, let $M \subseteq X \times Y \times Z$ be the set from $I_{3\mathrm{DM}}$, and let $X = \{s_1, \ldots, s_k\}$, $Y = \{s_{k+1}, \ldots, s_{2k}\}$, $Z = \{s_{2k+1}, \ldots, s_{3k}\}$. Similarly to the proof of theorem 4.8, we use Lemma 4.7 to construct an election with relative points as follows:

- $score(p) = 0$,
- for each $i \in \{1, \ldots, 3k\}$, $score(s_i) = k\alpha_1 - (k-1)\alpha_{m-3k+i} - \alpha_{1+r(i)}$,
- all dummy candidates have points such that adding at most k votes does not let them win the election.

For each $(x, y, z) = (s_h, s_i, s_j) \in M$ with $h < i < j$, we add an available voter voting as follows: $p > x > y > z > d_4 > \cdots > d_{m-3k-1} > s_1 > \cdots > s_{h-1} > d_1 > s_{h+1} > \cdots > s_{i-1} > d_2 > s_{i+1} > \cdots > s_{j-1} > d_3 > s_{j+1} > \ldots, s_{3k}$.

We say that x, y, and z are *approved* in this vote. Note that a candidate s_i gets $\alpha_{1+r(i)}$ (where $r(i) = 1$, 2 or 3 depending whether $s_i \in X$, Y or Z) points in a vote approving s_i, and gets α_{m-3k+i} in a vote not approving s_i.

We claim that p can be made a winner by adding at most k of the available votes if and only if the 3DM instance is positive.

First assume that the 3DM instance is positive, and let $C \subseteq M$ be a cover with $\|C\| = k$. We add the votes corresponding to the elements of C, i.e., for each candidate s_i, we add one vote that approves s_i and $(k-1)$ votes that do not approve s_i. The final score of the nondummy candidates is as follows:

- Candidate p gains α_1 points in each of the k added votes, so the final score of p is $k\alpha_1$,
- each candidate s_i gains $\alpha_{1+r(i)} + (k-1)\alpha_{m-3k+i}$ points from the one vote approving s_i and the $k-1$ votes not approving s_i. Thus the final score of s_i is $k\alpha_1 - (k-1)\alpha_{m-3k+i} - \alpha_{1+r(i)} + \alpha_{1+r(i)} + (k-1)\alpha_{m-3k+i} = k\alpha_1$.

Thus p and s_i tie, and so p is a winner of the election.

For the converse, assume that p can be made a winner by adding at most k voters. Let C be the elements of M corresponding to the added votes, then $\|C\| \leq k$.

Since $\alpha_1 \geq \alpha_4 > \alpha_{m-3k+1}$, we know that the score of s_1 before adding any votes is positive. Thus $C \neq \emptyset$, i.e., at least one vote is added. Let z be a candidate from Z who is approved in at least one of the added votes. We show that $\|C\| = k$. For this, indirectly assume that $\|C\| < k$. We show that z strictly beats p in this case: Consider the final scores of p and z. Let i be the index of z, i.e., the i with $s_i = z$. Note that $r(i) = 3$, since $s_i = z \in Z$.

- candidate p gains $\|C\|\alpha_1$ points,
- candidate z gains at least $\alpha_4 + (\|C\| - 1)\alpha_{m-3k+i}$ points (even more if z is approved in more than one of the additional votes). Thus the final score of z is at least

  $k\alpha_1 - (k-1)\alpha_{m-3k+i} - \alpha_4 + \alpha_4 + (\|C\| - 1)\alpha_{m-3k+i} = k\alpha_1 + (\|C\| - k)\alpha_{m-3k+i}$

To see that z strictly beats p, we compute the difference between their scores, which is

$$\begin{aligned} & \underbrace{k\alpha_1 + (\|C\| - k)\alpha_{m-3k+i}}_{score(s_i)} - \underbrace{\|C\|\alpha_1}_{score(p)} \\ = \; & (k - \|C\|)\alpha_1 - (k - \|C\|)\alpha_{m-3k+i} \\ = \; & (k - \|C\|)(\alpha_1 - \alpha_{m-3k+i}). \end{aligned}$$

Since, by assumption, $k - \|C\| > 0$, and $\alpha_1 \geq \alpha_4 > \alpha_{m-3k+1} \geq \alpha_{m-3k+i}$, this difference is positive and so s_i strictly beats p if $\|C\| < k$. Thus we know that $\|C\| = k$.

We now show that C is indeed a cover. Assume that this is not the case, so, since $\|C\| = k$, without loss of generality, there is some $z' \in Z$ which is covered twice, i.e., which is approved in at least two of the added votes. We show that this z' strictly beats p in the election after the addition of votes by comparing the score of p and z'. As above, let i be the index of z', i.e., chose i such that $z' = s_i$. The scores are as follows (again, $r(i) = 3$ since $s_i \in Z$):

- The candidate p again gets $k\alpha_1$ points, and so its final score is $k\alpha_1$.
- The candidate z' is approved in at least 2 of the k additional votes. Thus z' gains at least $2\alpha_4 + (k-2)\alpha_{m-3k+i}$ points. The final score of z' is thus at least

  $k\alpha_1 - (k-1)\alpha_{m-3k+i} - \underbrace{\alpha_{1+r(i)}}_{=\alpha_4} + 2\alpha_4 + (k-2)\alpha_{m-3k+i} = k\alpha_1 - \alpha_{m-3k+i} + \alpha_4.$

Since $\alpha_4 > \alpha_{m-3k+1} \geq \alpha_{m-3k+i}$, it follows that the final score of z' exceeds the score of p, and so z' indeed strictly beats p as claimed if C is not a cover. Thus C is indeed a cover as required. □

B.4 Proof of Theorem 4.10

Theorem 4.10. *f-CCAV is NP-complete for every* FP*-uniform pure $\mathbb{Q}$-generator f with $\alpha_4^{f,m} > \alpha_{m-2}^{f,m}$ for some m.*

Proof. Clearly, for pure scoring rules, if the condition $\alpha_4^{f,m} > \alpha_{m-2}^{f,m}$ is true for some m, then it remains true for all $m' \geq m$, since it is easy to see that $\alpha_4^{f,m+1} \geq \alpha_4^{f,m}$ and $\alpha_{m+1-2}^{f,m+1} \leq \alpha_{m-2}^{f,m}$.

We prove NP-hardness by a reduction from 3DM. So let $I_{3\mathrm{DM}}$ be a 3DM-instance, and let k be the cardinality of the set X in $I_{3\mathrm{DM}}$. Without loss of generality, we can assume that k is large enough such that the number $m = 6k$ satisfies the condition $\alpha_4^{f,m} > \alpha_{m-2}^{f,m}$. Due to the monotonicity of the coefficients and since $m - 3k + 1 = 3k + 1$, we have that

$$\alpha_4^{f,m} \geq \alpha_{m-3k+1}^{f,m} = \alpha_{3k+1}^{f,m} \geq \alpha_{m-2}^{f,m},$$

and since $\alpha_4^{f,m} > \alpha_{m-2}^{f,m}$, we know that one of the cases $\alpha_{3k+1}^{f,m} > \alpha_{m-2}^{f,m}$, or $\alpha_4^{f,m} > \alpha_{m-3k+1}^{f,m}$ occurs. In the first case, we use the reduction from Theorem 4.8, in the second case, the one from Theorem B.2. □

B.5 Proof of Theorem 4.12.1

Theorem 4.12. *The problem f-CCAV is NP-complete if f is one of the following pure generators:*

1. $f = (\alpha_1, \alpha_2, \alpha_3, \alpha_4, \ldots, \alpha_4, \alpha_5, 0)$ *with* $\alpha_2 > \alpha_4 > 0$.
2. $f = (\alpha_1, \alpha_2, \ldots, \alpha_2, 0)$ *with* $\alpha_1 \notin \{\alpha_2, 2\alpha_2\}$, $\alpha_2 > 0$.
3. $f = (\alpha_1, \alpha_2, \ldots, \alpha_2, \alpha_5, 0)$ *with* $\alpha_1 > \alpha_2 > \alpha_5$.
4. $f = (\alpha_1, \ldots, \alpha_1, \alpha_5, 0)$ *with* $\alpha_1 > \alpha_5 > 0$.

Proof. Again, we reduce from 3DM. Let $n = \|M\|$, and let $k = \|X\| = \|Y\| = \|Z\|$. We introduce the following candidates:

- The preferred candidate p,
- for each $c \in X \cup Y \cup Z$ a candidate c,
- two dummy candidates d_1 and d_2.

Using Lemma 4.7, we construct the registered voters such that the relative score of the candidates before adding voters is as follows:

- $score(p) = -k\alpha_4$,
- $score(d_1), score(d_2) < -k(\alpha_1 + \alpha_4)$,

- $score(x) = -\alpha_1 - (k-1)\alpha_4$ for each $x \in X$,
- $score(y) = -\alpha_2 - (k-1)\alpha_4$ for each $y \in Y$,
- $score(z) = -(k-1)\alpha_4$ for each $z \in Z$.

For each $(x, y, z) \in M$, we add one available voter voting

$$\underbrace{x}_{\alpha_1} > \underbrace{y}_{\alpha_2} > \underbrace{d_1}_{\alpha_3} > \underbrace{p > \ldots}_{\alpha_4} > \underbrace{d_2}_{\alpha_5} > \underbrace{z}_{0}.$$

We claim that p can be made a winner by adding at most k voters if and only if there is a cover $C \subseteq M$ with $\|I\| = k$.

First, assume that there is such a cover C. Then, adding the voters corresponding to the elements of C changes the scores as follows:

- p gains $k\alpha_4$ points and thus ends up with 0 points,
- each x (y) gains $\alpha_1 + (k-1)\alpha_4$ points ($\alpha_2 + (k-1)\alpha_4$ points), which leads to 0 points,
- each z gains $(k-1)\alpha_4$ points, also finishing with 0 points,
- by construction, d_1 and d_2 cannot win the election.

For the other direction, assume that C is a set of added voters that makes p win the election with $\|C\| \leq k$. Since $\alpha_4 > 0$, we know that p must gain one point against every candidate $z \in Z$, and can gain against one of these candidates for each added vote, at least $\|Z\| = k$ votes must be added, so $\|C\| = k$. We prove that C is a cover. By the above, C covers Z. Indirectly assume that there is some x (y) that is voted in the first (second) place in more than one of the voters from C. In this case, x (y) gains at least $2\alpha_1 + (k-2)\alpha_4$ ($2\alpha_2 + (k-2)\alpha_4$) points, and so ends up with $-\alpha_1 - (k-1)\alpha_4 + 2\alpha_1 + (k-2)\alpha_4 = \alpha_1 - \alpha_4$ points ($\alpha_2 - \alpha_4$ points). Since $\alpha_1 \geq \alpha_2 > \alpha_4$, p does not win the election in this case, a contradiction. This completes the proof. □

B.6 Proof of Theorem 4.12.2

Theorem 4.12. *The problem f-CCAV is NP-complete if f is one of the following pure generators:*

1. $f = (\alpha_1, \alpha_2, \alpha_3, \alpha_4, \ldots, \alpha_4, \alpha_5, 0)$ *with* $\alpha_2 > \alpha_4 > 0$.
2. $f = (\alpha_1, \alpha_2, \ldots, \alpha_2, 0)$ *with* $\alpha_1 \notin \{\alpha_2, 2\alpha_2\}$, $\alpha_2 > 0$.
3. $f = (\alpha_1, \alpha_2, \ldots, \alpha_2, \alpha_5, 0)$ *with* $\alpha_1 > \alpha_2 > \alpha_5$.
4. $f = (\alpha_1, \ldots, \alpha_1, \alpha_5, 0)$ *with* $\alpha_1 > \alpha_5 > 0$.

The proof of this theorem uses two different reductions depending on which of the cases $\alpha_1 > 2\alpha_2$ or $\alpha_1 < 2\alpha$ applies. We start with the case $\alpha_1 > 2\alpha_2$:

Theorem B.3. *f-CCAV is NP-complete for the generator $f = (\alpha_1, \alpha_2, \ldots, \alpha_2, 0)$ with $\alpha_1 > 2\alpha_2 > 0$.*

Proof. We again reduce from 3DM. So let $k = \|X\| = \|Y\| = \|Z\|$. We introduce the following candidates:

- the preferred candidate p,
- a candidate c for each $c \in X \cup Y \cup Z$,
- for each $S_i \in M$, a candidate S_i,
- a dummy candidate d (used for the application of Lemma 4.7).

Without loss of generality, due to Proposition 2.1, we can assume that for the relevant length, $\gcd \alpha_1, \alpha_2 = 1$, and $\alpha_1, \alpha_2 \in \mathbb{N}$. In particular, the number 1 can be obtained as a linear combination of α_1 and α_2. We thus can apply Lemma 4.7 to construct a set of voters V such that the relative scores are as follows:

- $score(p) = 0$,
- $score(y) = \alpha_2$ for each $y \in Y$,
- $score(x) = score(z) = \alpha_2 - \alpha_1$ for each $x \in X$ and $z \in Z$,
- $score(S_i) = 2\alpha_2 - \alpha_1 + 1$ for each $S_i \in M$,
- $score(d) < -3k\alpha_1$.

For each $S_i = (x_i, y_i, z_i) \in M$, we add the following available voters:

- $S_i > \cdots > y_i$,
- $z_i > \cdots > S_i$,
- $x_i > \cdots > S_i$.

We say that a vote *vetoes* the candidate it places in the last position. We claim that p can be made a winner of the election with adding at most $3k$ votes if and only if the 3DM-instance is positive. By construction, the dummy candidate d cannot win the election by adding at most $3k$ votes, so we ignore d in the sequel.

Note that since p is voted in one of the middle positions in all of the potentially added voters, the score of p only depends on the number of added voters, and not on the concrete choice of added votes. We thus consider only the change of relative scores between the candidates and p when studying the effect of added voters.

First assume that there is a cover C with $\|C\| = k$. For all $S_i \in C$, we add all votes voting S_i in the first or in the last position. Then the scores (relatively to p) change as follows:

- each y_i is vetoed once, and so loses α_2 points relative to p, so has final score 0,
- each S_i with $i \in I$ appears once in the first position and twice in the last position, so S_i gains $\alpha_1 - \alpha_2 - 2\alpha_2 = \alpha_1 - 3\alpha_2$ points. So the final score of such an S_i is $2\alpha_2 - \alpha_1 + 1 + \alpha_1 - 3\alpha_2 = -\alpha_2 + 1 \leq 0$, since $\alpha_2 \geq 1$.
- for each S_i with $i \notin I$, the relative points of S_i and p do not change. Since $\alpha_1 > 2\alpha_2$ and $\alpha_1, \alpha_2 \in \mathbb{N}$, it follows that $2\alpha_2 - \alpha_1 + 1 \leq 0$ and so S_i does not beat p.
- each z_i and each x_i gain exactly $\alpha_1 - \alpha_2$ points against p, and so tie with p.

Thus p wins the election after adding these voters.

For the converse, assume that there is a set V of available voters that we can add with $\|V\| \leq 3k$ such that p wins the resulting election. Since each $y \in Y$ is currently winning against p, and the only way for y to lose points relatively to p is adding the vote voting y last, we know that for each $y \in Y$, there is one vote in V that votes y last. In particular, we have $\|V\| \geq k$. So there are votes $v_1, \ldots, v_k \in V$, each voting a different y in the last position, and thus each voting a different S_i in the first position.

For each such vote, one of the S_i gains $(\alpha_1 - \alpha_2)$ points against p. If such an S_i is vetoed only once, then its final score is $2\alpha_2 - \alpha_1 + 1 + \alpha_1 - \alpha_2 - \alpha_2 = 1$, and so p does not win the election. Thus each S_i gaining points in one of the votes $v_1, \ldots, v_k$ must be vetoed at least twice. Since $\|V\| \leq 3k$, we know that $\|V\| = 3k$. Let $I = \{i \mid S_i > \cdots > y_i \in V\}$. We claim that I is a cover. By the above, I covers each $y \in Y$. Now assume that for $i \neq j \in I$, we have $x_i = x_j$. Then x_i gains $2(\alpha_1 - \alpha_2)$ points against p, and so p does not win the election, since $\alpha_1 > \alpha_2$. The same argument holds for y. Thus, due to cardinality, we have a cover of X, Y, and Z. □

The case $\alpha_1 < 2\alpha_2$ is similar:

Theorem B.4. *f-CCAV is NP-complete for the generator $f = (\alpha_1, \alpha_2, \ldots, \alpha_2, 0)$ with $\alpha_1 < 2\alpha_2$.*

Proof. Very similar to the proof of Theorem B.3. We again reduce from 3DM. We use the same candidate set as in the proof of Theorem B.3 (including the dummy candidate for the application of Lemma 4.7), and set up the scores of the nondummy candidates as follows:

- $score(p) = 0$,
- $score(y) = \alpha_2 - \alpha_1$ for each $y \in Y$,
- $score(x) = score(z) = \alpha_2$ for each $x \in X$ and each $z \in Z$,
- $score(S_i) = min(0, 3\alpha_2 - 2\alpha_1)$.

For each tuple $S_i = (x_i, y_i, z_i)$, we add three available voters

- $y_i > \cdots > S_i$,

- $S_i > \cdots > x_i$,
- $S_i > \cdots > z_i$.

We say that a vote as above *approves* (*vetoes*) the candidate put in its first (last) position.

If there is a cover with size at most k, we choose the 3 voters associated with each element of the cover. This lets each x and each z lose α_2 points against p, each y_i gains $(\alpha_1 - \alpha_2)$ points against p, and each S_i in the cover ends up with $score(S_i) + 2(\alpha_1 - \alpha_2) - \alpha_2 = score(S_i) + 2\alpha_1 - 3\alpha_2 \leq 3\alpha_2 - 2\alpha_1 + 2\alpha_1 - 3\alpha_2 = 0$ points (all points counted relative to p). Thus all candidates tie and p is a winner of the resulting election.

For the converse, assume that p can be made a winner by adding at most $3k$ voters, let V be the corresponding set. Clearly, each $x \in X$ and each $z \in Z$ need to be vetoed by some vote in V. Thus there are at least $2k$ votes in V that vote some S_i in the first position. Let C contain all indices S_i such that there is a vote in V approving of S_i. Clearly, $\|C\| \geq k$. Then, each $S_i \in C$ gains at least $(\alpha_1 - \alpha_2)$ points from these votes, and thus has $min(0, 3\alpha_2 - 2\alpha_1) + \alpha_1 - \alpha_2$ points. If the minimum is 0, then this clearly is more than 0, if the minimum is $3\alpha_2 - 2\alpha_1$, then the score adds up to $3\alpha_2 - 2\alpha_1 + \alpha_1 - \alpha_2 = 2\alpha_2 - 1\alpha_1$, which also exceeds 0 since $\alpha_1 < 2\alpha_2$. Thus each $S_i \in C$ must be vetoed at least once. In particular, there must be at least k votes vetoing some S_i, and since due to cardinality reasons, there can be only k votes vetoing some S_i, we know that $\|C\| \leq k$. Due to the above, we also know $\|C\| \geq k$, so $\|C\| = k$.

It remains to show that C covers X, Y, and Z. As argued above, each x and each z need to be vetoed once, so C covers X and Z. To show that C also covers Y, it suffices to show that no y can be approved twice. This trivially follows since $score(y) = \alpha_2 - \alpha_1$, and each approval lets y gain $\alpha_1 - \alpha_2 > 0$ points against p. □

The proof of Theorem 4.12.2 now directly follows from the above results:

Proof. If $\alpha_1 \geq \alpha_2 > 0$ with $\alpha_1 \notin \{\alpha_2, 2\alpha_2\}$, then one of the following cases applies:

- $\alpha_1 > 2\alpha_2 > 0$, in this case the result follows from Theorem B.3,
- $2\alpha_2 > \alpha_1 > \alpha_2 > 0$, in this case the result follows from Theorem B.4. □

B.7 Proof of Theorem 4.12.3

Theorem 4.12. *The problem f-CCAV is NP-complete if f is one of the following pure generators:*

1. $f = (\alpha_1, \alpha_2, \alpha_3, \alpha_4, \ldots, \alpha_4, \alpha_5, 0)$ *with* $\alpha_2 > \alpha_4 > 0$.
2. $f = (\alpha_1, \alpha_2, \ldots, \alpha_2, 0)$ *with* $\alpha_1 \notin \{\alpha_2, 2\alpha_2\}$, $\alpha_2 > 0$.
3. $f = (\alpha_1, \alpha_2, \ldots, \alpha_2, \alpha_5, 0)$ *with* $\alpha_1 > \alpha_2 > \alpha_5$.

4. $f = (\alpha_1, \ldots, \alpha_1, \alpha_5, 0)$ *with* $\alpha_1 > \alpha_5 > 0$.

Proof. We again reduce from 3DM, let $k = \|X\| = \|Y\| = \|Z\|$. Using Lemma 4.7, we set up the the relative scores as follows:

- $score(p) = 0$,
- $score(x) = -(\alpha_1 - \alpha_2)$ for each $x \in X$,
- $score(y) = \alpha_2 - \alpha_5$ for each $y \in Y$,
- $score(z) = \alpha_2$ for each $z \in Z$,
- $score(d) < -k\alpha_1$, here d again is a dummy candidate required for the application of Lemma 4.7, who cannot win the election and whom we ignore in the sequel.

The available voters are as follows: For each (x, y, z) in M, there is an available voter voting

$$x > \cdots > y > z.$$

We claim that p can be made a winner of the election if and only if the 3DM-instance is positive. First assume that the instance is positive, we then add the k votes corresponding to the cover. Then:

- each $x \in X$ gains exactly $\alpha_1 - \alpha_2$ points relatively to p,
- each $y \in Y$ loses exactly $\alpha_2 - \alpha_5$ points relatively to p,
- each $z \in Z$ loses exactly α_2 points relatively to p.

Thus all candidates tie and p is a winner of the election. For the converse, assume that p can be made a winner by adding at most k voters, let C be the corresponding subset of M. Since each $y \in Y$ and each $z \in Z$ need to lose points relatively to p, C covers each Y and Z. C cannot cover any $x \in X$ twice, since otherwise x would gain $2(\alpha_1 - \alpha_2)$ points and thus beat p in the election. Thus C covers each x at most once, and so covers each x exactly once, thus C is a cover as required. □

B.8 Proof of Theorem 4.12.4

Theorem 4.12. *The problem f-CCAV is NP-complete if f is one of the following pure generators:*

1. $f = (\alpha_1, \alpha_2, \alpha_3, \alpha_4, \ldots, \alpha_4, \alpha_5, 0)$ *with* $\alpha_2 > \alpha_4 > 0$.

2. $f = (\alpha_1, \alpha_2, \ldots, \alpha_2, 0)$ *with* $\alpha_1 \notin \{\alpha_2, 2\alpha_2\}$, $\alpha_2 > 0$.

3. $f = (\alpha_1, \alpha_2, \ldots, \alpha_2, \alpha_5, 0)$ *with* $\alpha_1 > \alpha_2 > \alpha_5$.

4. $f = (\alpha_1, \ldots, \alpha_1, \alpha_5, 0)$ *with* $\alpha_1 > \alpha_5 > 0$.

Proof. We again reduce from 3DM. Let M be a 3DM-instance with $\|M\| = n$, and let $k = \|X\| = \|Y\| = \|Z\|$. In addition to the preferred candidate p, we introduce a candidate c for each $c \in X \cup Y \cup Z$, and for each $S_i \in M$, we introduce a candidate S_i and a candidate S_i'. We use Lemma 4.7 to set up the registered voters such that, relatively to p, we have the following scores:

- $score(p) = 0$,
- $score(x) = score(y) = score(z) = \alpha_1$ for each $x \in X$, $y \in Y$, and $z \in Z$,
- $score(S_i) = min(\alpha_1, 2(\alpha_1 - \alpha_5))$ for each $S_i \in M$,
- $score(S_i') = \alpha_1 - \alpha_5$ for each $S_i \in M$,
- $score(d) < -(n+2k)\alpha_1$, where d again is a dummy candidate needed to apply Lemma 4.7, whom we ignore from now on.

Note that since $\alpha_1 > \alpha_5 > 0$, all of these scores—except for p and d—are strictly positive. For each $S_i = (x_i, y_i, z_i) \in M$, we add the following available voters:

- $\cdots > S_i' > S_i$,
- $\cdots > S_i > x_i$,
- $\cdots > S_i > y_i$,
- $\cdots > S_i' > z_i$.

Again, we say that a vote *vetoes* its last-places candidate. We claim that p can be made a winner by adding at most $n + 2k$ voters if and only if the 3DM-instance is positive. First assume that there is a cover $C \subseteq M$ with $\|C\| \leq k$. We add the following voters:

- for each $S_i = (x_i, y_i, z_i) \in C$, we add the voters $\cdots > S_i > x_i$, $\cdots > S_i > y_i$, and $\ldots S_i' > z_i$.
- for each $S_i \notin C$, we add the voter $\cdots > S_i' > S_i$.

When adding these voters, each candidate $c \in X \cup Y \cup Z$ loses α_1 points against p and thus ends with 0 points. Each candidate S_i' loses $\alpha_1 - \alpha_5$ points and so also has 0 points in the end. A candidate $S_i \in C$ loses $2(\alpha_1 - \alpha_5) \geq score(S_i)$ points, and a candidate $S_i \notin C$ loses $\alpha_1 \geq score(S_i)$ points. In both cases, S_i has at most 0 points after the election. Thus p wins the election.

For the converse, assume that there is a set V of voters with $\|V\| \leq n + 2k$ such that p wins the election after the votes in V are added. Without loss of generality, we can assume $\|V\| = n + 2k$, since none of the available votes hurt p. Let I be the set of indices i such that a vote of the form $\cdots > S_i > x_i$, $\cdots > S_i > y_i$, or $\cdots > S_i' > z_i$ is contained in V. We observe the following:

- Since each candidate from X, Y, and Z must lose α_1 points against p, for each of these candidates there must be a vote in V voting that candidate last. Thus at least $3k$ votes voting one such candidate last must appear in V.
- Without loss of generality, we can assume that exactly $3k$ votes of the above form appear in V, since it does not help to veto one of these candidates twice (to stop one of the candidates S_i, S'_i from winning against p, a vote of the form $\cdots > S'_i > S_i$ is always preferable).
- Thus there are exactly $n - k$ votes of the form $\cdots > S'_i > S_i$ in V. Since one such vote is enough to lower the relative score of S'_i and S_i to the score of p and below, we can without loss of generality assume that these votes are all distinct.
- Let C be the set of indices such that $\cdots > S'_i > S_i \notin V$. Due to the above, we know that $I = k$.
- For each $i \in C$, since S_i must lose more than $\alpha_1 - \alpha_5$ points, both voters $\cdots > S_i > x_i$ and $\cdots > S_i > y_i$ must appear in V.
- For each $i \in C$, since S'_i must lose points, the votes $\cdots > S'_i > z_i$ must appear in V,
- Since there are only $3k$ voters voting a candidate from $X \cup Y \cup Z$ last, it follows that C is a cover. □